EDN®
VOICE OF THE ENGINEER
MAR 17
Issue 6/2011
www.edn.com

COMPLETE IC SIMULATION

REQUIRES A FULL TOOLBOX OF HARDWARE AND SOFTWARE

Page 38

SURPASSING THE BANDWIDTH LIMITATIONS OF CACHE-BASED PROCESSING ARCHITECTURES

Page 30

HARDWARE CHOICES FOR MPLS AND BACKBONE BRIDGING

page 50

Visit digikey.com for the industry's broadest selection of semiconductors!

AUTHORIZED DISTRIBUTOR FOR MORE THAN 440 SUPPLIER PARTNERS.*

The industry's broadest product selection available for immediate delivery

www.digikey.com

1.800.344.4539

*Digi-Key is an authorized distributor for all supplier partners. New products added daily.

NEDA Supplier Authorized Distributors

The World's Standard.

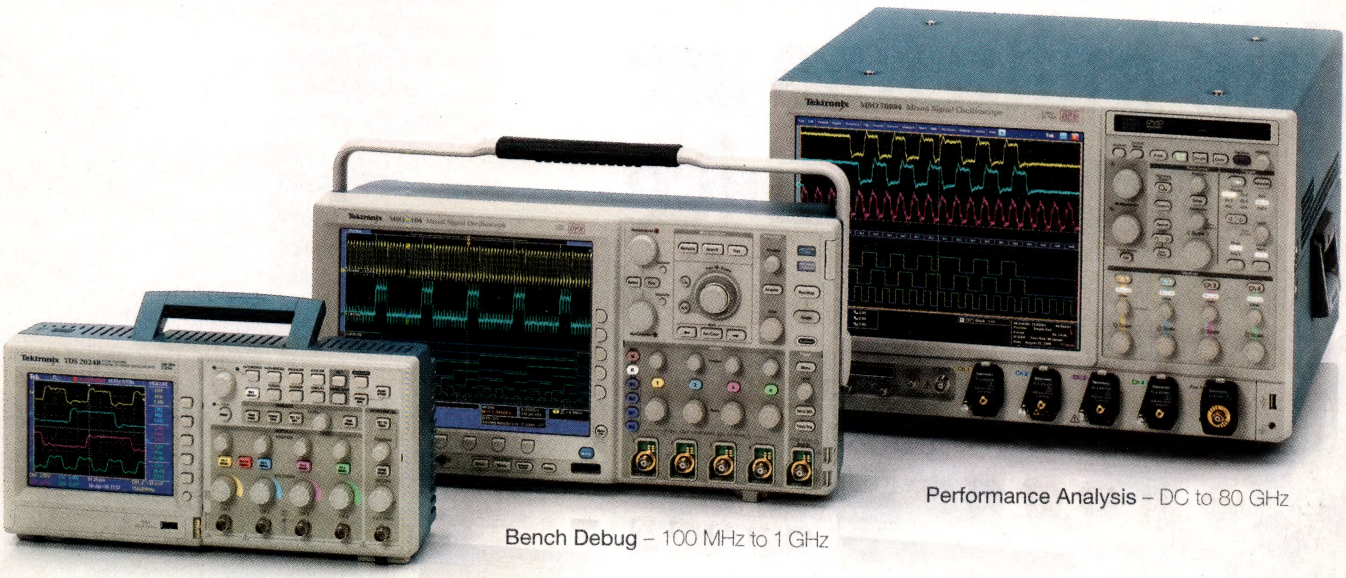

Basic Visualization – 40 to 500 MHz

Bench Debug – 100 MHz to 1 GHz

Performance Analysis – DC to 80 GHz

Oscilloscopes from Basic to Bench and Beyond.

Whatever your need, there is a Tektronix oscilloscope to fill it. Tektronix offers the world's broadest portfolio of oscilloscopes to debug and test tomorrow's designs, today. Start with Tektronix Basic oscilloscopes: fast, familiar and affordable; they're what you know and trust. Our Bench oscilloscopes provide next-level productivity, with the feature-rich tools you need to debug today's complex mixed signal designs. And for those who demand fast, flexible, in-depth PHY layer analysis, our Performance oscilloscopes offer the industry's best signal fidelity, verification and characterization capabilities to help you shorten your design cycles.

No wonder **8 of 10 engineers worldwide trust Tektronix** to help them bring advanced designs to market on time and on budget. Wherever you're going, we'll get you there.

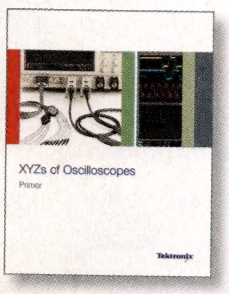

See why Tektronix is the World's Standard in Oscilloscopes and download your free Oscilloscope Primer.

www.tektronix.com/worldstandard

Tektronix®

Why Coilcraft wirewound chip inductors are your #1 choice

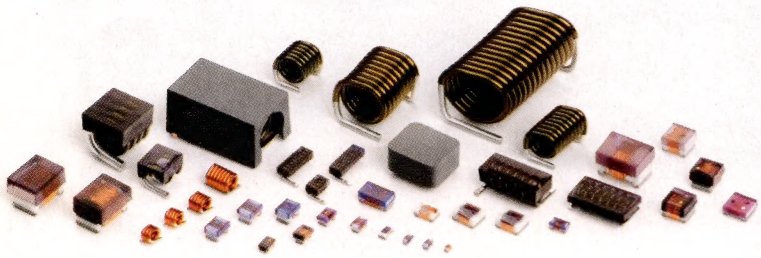

Higher Q Compared to non-wirewounds, our chip inductors usually have Qs that are 50 to 150% higher.

Lower DCR Put up to 3 times more current through our chip inductors thanks to their low DC resistance.

Higher SRF The solenoid winding of our inductors gives them a much higher SRF than multilayer parts.

Tighter tolerance Precision manufacturing lets us consistently make parts with ±2% inductance tolerance. Many popular values also come in ±1%.

Better support With our engineer-friendly web site, interactive design tools and generous free samples, Coilcraft is just plain easier to do business with.

Visit **www.coilcraft.com** for information on all our high performance wirewound inductors.

coilcraftdirect.com
No min order. Next day delivery.

Coilcraft®

WWW.COILCRAFT.COM

EDN 3.17.11 contents

Complete IC simulation requires a full toolbox of hardware and software

by Mike Demler, Technical Editor

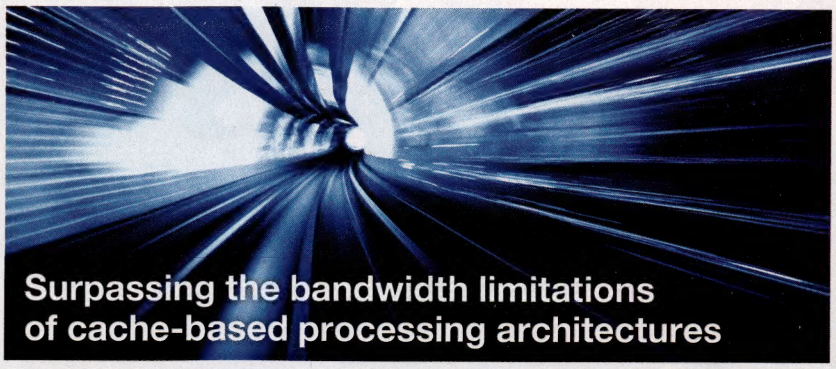

Surpassing the bandwidth limitations of cache-based processing architectures

by Daniel Proch, Netronome

pulse

COVER: PHOTO-ILLUSTRATION BY TIM BURNS. TOOLBOX: SERGEY KULIKOV/ISTOCKPHOTO.COM; CHIPS: CRAFTVISION/ISTOCKPHOTO.COM; PROBES: IGOR SMICHKOV/ISTOCKPHOTO.COM

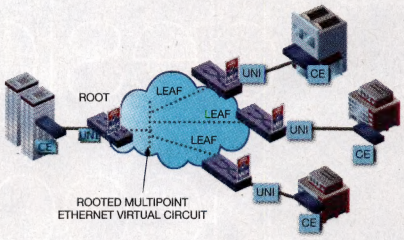

Hardware choices for MPLS and backbone bridging

by Martin Nuss, PhD, Vitesse

DESIGNIDEAS

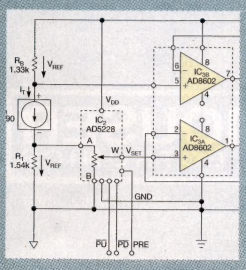

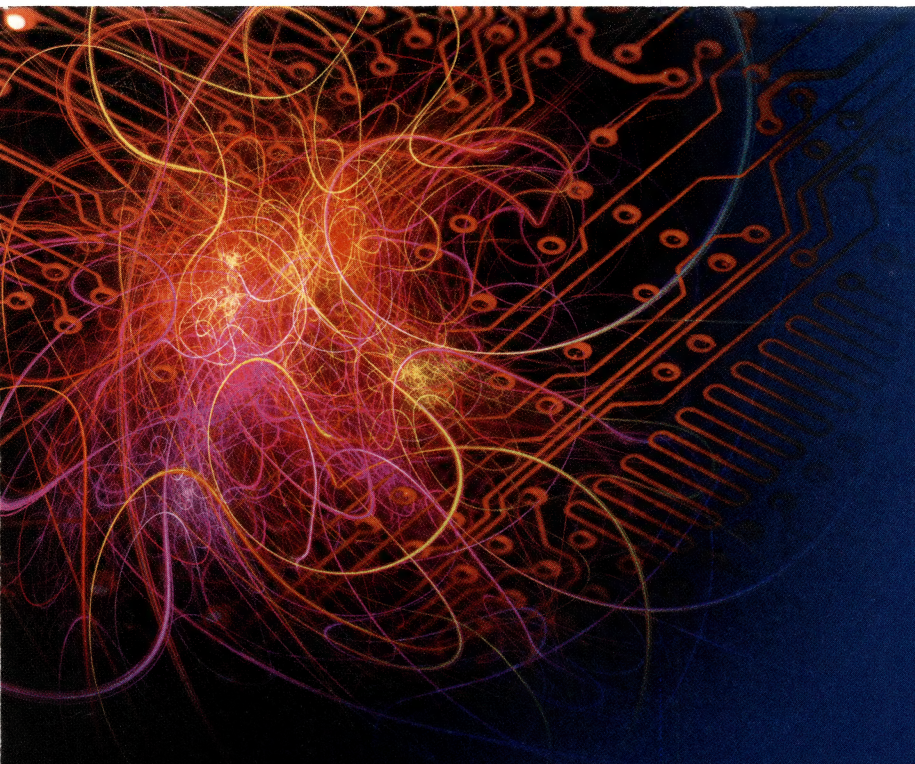

mouser.com
Semiconductors and electronic components for design engineers.

We've got what design engineers need. **The future.**

You'll find the newest products for your newest designs.
Get what's next at Mouser.

Scan Here

mouser.com

MOUSER ELECTRONICS
a tti company

contents 3.17.11

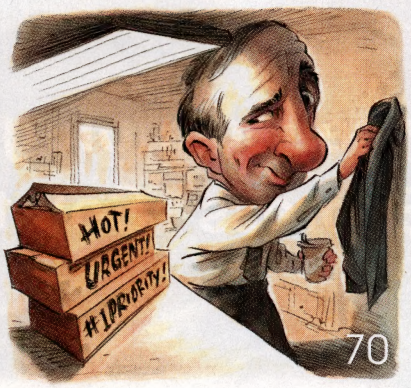

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out this Web-exclusive article:

Another round in the patent wars
Two of three Intellectual Ventures lawsuits target semiconductor companies.
→www.edn.com/110317toca

iGen: the next generation of engineers

Three years ago in the face of declining interest among students in the United States in engineering-degree programs, *EDN* sister publication *EE Times* set out to build an online space aimed at K-12 students that sparked and reinforced a passion for math and science, created rock stars of engineers, and showed kids that the stuff that engineers and scientists do is really fun and cool. iGen (Innovation Generation) is the online place for kids who have a passion for math, science, and technology to publish articles as student reporters, participate in basic design challenges, share their projects, and explore all the cool things that scientists and engineers create. Get inspired by their excitement and their designs on this site.
→www.edn.com/110317tocb

EDN® (ISSN#0012-7515) is published semimonthly, 24 times per year, by UBM Electronics, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 310/445-4200; FAX 310/445-4299. Periodicals postage paid at Los Angeles, CA, and at additional mailing offices. SUBSCRIPTIONS—Free to qualified subscribers as defined on the subscription card. Rates for nonqualified subscriptions, including all issues: US, $150 one year; $250 two years; $300 three years. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. For telephone inquiries regarding subscriptions, call 763/746-2792. E-mail: EDN@kmpsgroup.com. CHANGE OF ADDRESS—Notices should be sent promptly to PO Box 47461, Plymouth, MN 55447. Please provide old mailing label as well as new address. Allow two months for change. NOTICE—Every precaution is taken to ensure accuracy of content; however, the publisher cannot accept responsibility for the correctness of the information supplied or advertised or for any opinion expressed herein. POSTMASTER—Send address changes to EDN, PO Box 47461, Plymouth, MN 55447. Canada Post: Publications Mail Agreement 40612608. Return undeliverable Canadian addresses to Pitney Bowes Inc, PO Box 25542, London, ON N6C 6B2.

Volume 56, Number 6 (Printed in USA).

Integrated Device Technology – Leader in PCI Express® Solutions

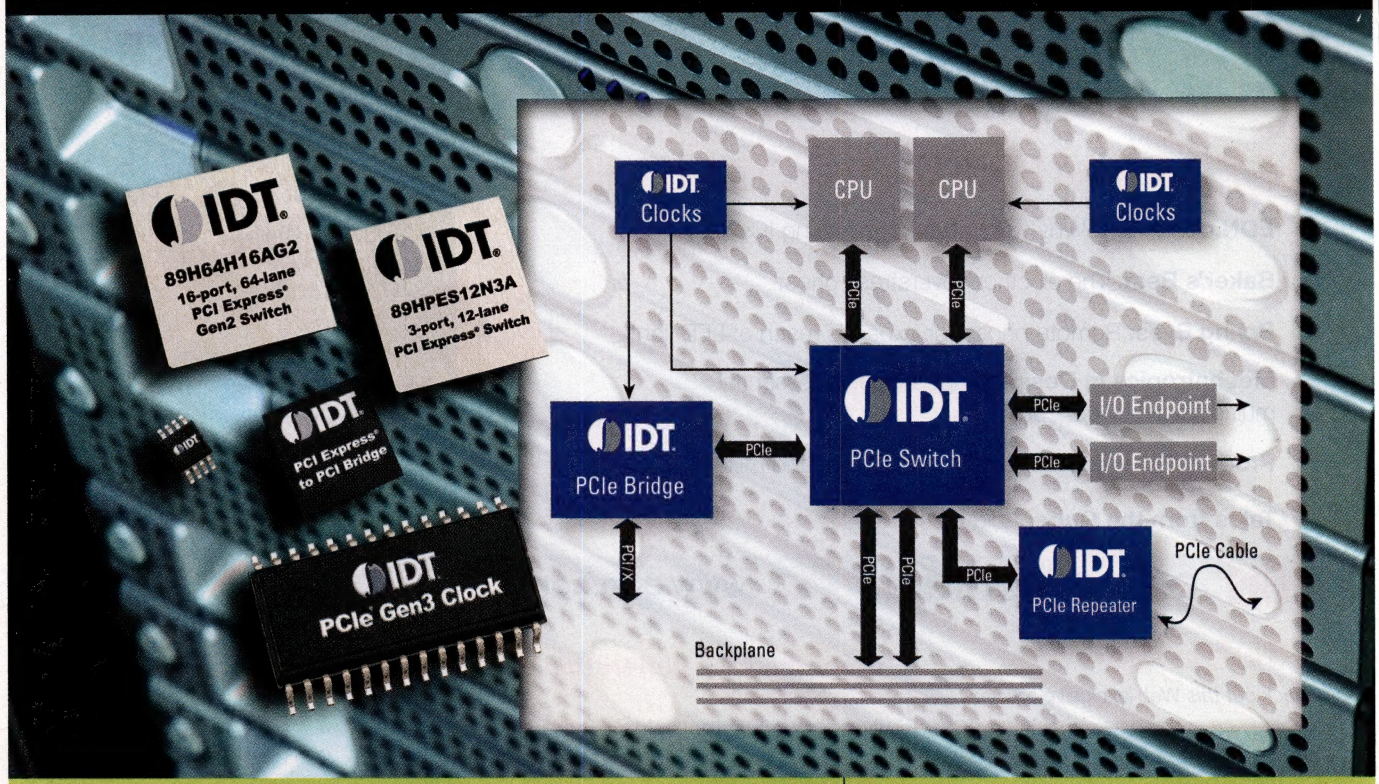

PCI Express Switches
Industry's most comprehensive family of high-performance, scalable PCIe switching solutions
- Gen 1 and Gen 2 Portfolio
- Up to 64 lane and 24 port devices
- Highly flexible port configurations
- Unprecedented 8 non-transparent bridging (NTB) functions to enable multi-root applications

PCI Express Repeaters
5Gbps Gen 2 PCIe Signal Repeaters
- Four and eight channel signal repeaters
- Compensates for cable and PCB trace attenuation and ISI jitter
- Programmable receiver equalization to 30dB and de-emphasis to -8.5dB
- Leading edge power minimization in active and shutdown modes
- No external bias resistors or reference clocks required
- Channel mux mode, demux mode, 1 to 2 channels multicast, and Z-switch function mode

PCI Express Bridges
High performance PCIe bridging to legacy PCI and PCI-X protocols
- Ultra low power version for consumer applications
- Forward mode buffer optimization
- Only PCI Express bridges with Short Term Caching for significant improvement in PCI Read performance
- Pin compatibility with competitive offering for dual sourcing solution

PCI Express Clocking Solutions
Industry's broadest offering of PCIe Gen1, Gen2 and Gen3 clock generation and buffering solutions
- World's First PCI Express Gen3 Family of Timing Devices
- Clock Generators and Synthesizers
- Spread Spectrum Clock Generators
- Fanout Buffers and Muxes
- Jitter Attenuators
- Frequency Translators
- Zero Delay Buffers

IDT
Integrated Device Technology

The Analog and Digital Company™

For more information visit us online at www.IDT.com

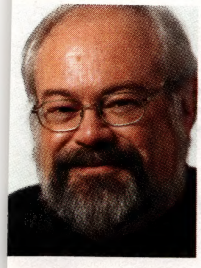

BY RON WILSON, EDITORIAL DIRECTOR

Why write about networking ICs?

It's a reasonable question. *EDN* is for a general audience of electronics-design engineers. Networking is a specific application, employing a minority of *EDN*'s readers and, one could argue, increasingly isolated from the mainstream of electronics technology. Where else is 10 Gbps routine? Who else uses content-addressable memory?

Specifically because of this divergence, it is vital to understand what is going on in the networking world. I would like to approach this argument from several points of view: underlying technology, implementation, system architecture, and our very way of thinking about systems.

First, what's causing networking to move away from the mainstream? In a word: bandwidth. Our insatiable demand for diversion by slap-in-the-face media experiences is driving us to consume unthinkable amounts of bandwidth, nearly all the time and everywhere. That appetite forces service providers all through the network to move bits faster and more efficiently. Faster means higher I/O frequencies. Efficiently means inspecting the insides of individual packets, classifying them, and prioritizing their transmission. The two forces work against each other. As the packet rate increases, the amount of processing per packet is increasing even faster. These pressures have forced networking engineers to solve problems unique—at least in scale—to their field. In doing so, they have blazed trails the rest of us may well follow.

Let's begin with the obvious. If you are to handle 10- or 100-Gbps packet streams, you are going to learn a lot about working with signals faster than 1 GHz—often far faster. Accordingly, the networking designers have given us affordable FPGAs with serial I/O pins capable of multigigahertz transmission, interconnect that can move such signals across a backplane, and board-level signal-integrity tools that can analyze such links. The availability of these capabilities is beginning slowly to make trans-gigahertz serial I/O an alternative to wide parallel buses in non-networking applications. In some designs, fast serial links routinely serve as chip-to-chip interconnect.

At a higher level of abstraction, consider packet-processing engines. To keep up with demand, networking designers have not only wrung every drop from Moore's Law but also pioneered the exploration of real-time multiprocessing. They have tried just about every combination of parallel and pipelined organization, hardware and software, accelerators, and multicore clusters. They have learned about the relationships between task mix, data profile, microarchitecture, throughput, and energy use. These lessons will be necessary to the rest of us as we confront increasing data loads and algorithm complexity, in the shadow of severe energy constraints.

At the architectural level, network-switch and router developers have pressed many new concepts into service. They use intelligent memory structures. They pioneered the separation of control and data planes in an application. They take on the challenges of security, reliability, and self-repair in large-scale systems. All this learning is likely to come in handy for the rest of us as we face increasingly demanding designs and uncertain process technology.

But the changes born in the networking world may include more fundamental ideas. Long ago, moving information between nodes meant modulating analog signals. Decades ago, we began the transition to discontinuous modulation; most signals became digital. Analog designers saw digital as an unconscionable waste of bandwidth and energy.

The architects and designers in networking are, as they struggle with their own problems, writing chapters of our future.

We may now be on the threshold of another transition, to a world in which most information moves between nodes not in raw digital form but as Internet Protocol packets. To a digital-system designer, it sounds hideously wasteful, but we have no idea what architectural vistas this concept might open to us.

So why write about complex, highly specialized designs for network packet processing? It's not just that the designs are interesting attempts to attack leading-edge challenges. It's that the architects and designers in networking are, as they struggle with their own problems, writing chapters of our future. EDN

Contact me at ron.wilson@ubm.com.

EYEIDEA/ISTOCKPHOTO.COM

Support **Across The Board.**™
From Design to Delivery™

Now, you can have it all.™

Faster and easier than ever before. Our commitment to customer service is backed by an extensive product offering combined with our supply chain and design chain services – which can swiftly be tailored to meet your exact needs. We have dedicated employees who have the experience to provide the highest level of customer service with accuracy and efficiency. All of our technical experts are factory certified on the latest technologies, providing you the expertise to move projects forward with speed and confidence.

Avnet offers the best of both worlds: extensive product and supply chain knowledge, and specialized technical skill which translates into faster time to market – and the peace of mind that comes from working with the industry's best. **Avnet is ranked Best-In-Class* for well-informed sales reps, knowledgeable application engineers and our design engineering services – proof that we consistently deliver:**

> Industry recognized product expertise
> Specialized technical skills

Ready. Set. Go to Market.™
Visit the Avnet Design Resource Center™ at:
www.em.avnet.com/drc

SAME-DAY SHIPPING
orders placed before 8:00AM. central time

AVNET®
electronics marketing

Accelerating Your Success™

50 Years AVNET ON THE NEW YORK STOCK EXCHANGE NYSE.

Supplier Authorized Distributor

1 800 332 8638
www.em.avnet.com

Follow us on Twitter!
www.twitter.com/avnetdesignwire

*As rated by Hearst Electronics Group: The Engineer & Supplier Interface Study, 2009.

EDN

UBM Electronics

EDN, 33 Hayden Avenue, Lexington, MA 02421 **www.edn.com. Subscription inquiries:** 1-763-746-2792; EDN@kmpsgroup.com. **Address changes:** Send notice promptly to PO Box 47461, Plymouth, MN 55447. Please provide an old mailing label as well as your new address. Allow two months for the change.

UBM Electronics, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 1-310-445-4200; fax: 1-310-445-4299.

EDITED BY FRAN GRANVILLE

pulse

INNOVATIONS & INNOVATORS

Multicore embedded processors accelerate 4G LTE applications

ARM recently introduced the Cortex-R5 and Cortex-R7 embedded real-time processors targeting LTE (long-term-evolution) and LTE-Advanced applications. The new designs build on the capabilities of the ARM Cortex-R4 processor, which the company introduced in 2006, by maintaining binary compatibility.

The Cortex-R complements specialized DSP processors in 4G (fourth-generation)-modem applications. You can use the R5 and the R7 in single- or dual-core configurations. A dual-core configuration reduces software overhead by allowing simultaneous processing of the download and upload channels that an LTE paired-spectrum FDD (frequency-division-duplexing) wireless-network architecture uses. ARM will manufacture the CPU IP (intellectual-property) cores in a low-power, 28-nm CMOS process.

The R5 adds to the features of the previous-generation ARM Cortex-R4 by extending error management to the buses, increasing transaction reliability. You can configure the device during synthesis to optimize power efficiency. For example, you can configure the R5 with various cache sizes, and you can use single- or double-precision FPUs (floating-point units) with the licensed IP to save real estate. Other features of the R5 include a high-priority low-latency peripheral port, and an accelerator-coherency port provides cache coherency for increased data-transfer efficiency and more reliable firmware.

The Cortex-R7 is 50% more efficient than the R5 and adds SMP (symmetrical multiprocessing) with coherent caches between processors. The R7 supports full-fledged SMP without the need for an MMU (memory-management unit), limited by the ability of the 32-bit addressing to access as much as 4 Gbits of memory.

The R7 enables you to extend error recovery to hard errors, such as stuck-at faults of as many as four errors per memory bank, configurable through software for application-defined policies. Both the R5 and R7 support lock-step operation for failure backup in one CPU. The Cortex-R7 allows out-of-order instruction issue and completion. An R7 register-renaming feature lets you maximize processing resources.

—by Mike Demler

ARM, www.arm.com.

TALKBACK

"The story is about the Czar of Russia and the quiet, compressed-air-driven starter of his Delaunay-Belleville. In 1911, it became sold as an option."

—Physicist and *EDN* reader "MorDuGuz," in *EDN*'s Talkback section, at http://bit.ly/dFIkfr. Add your comments.

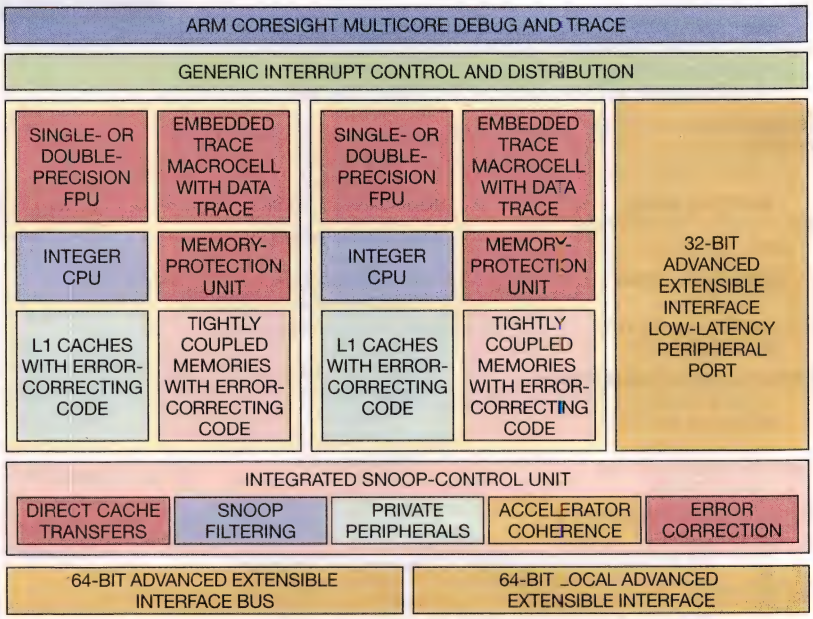

You can use the ARM Cortex-R7 real-time processor in a single- or dual-core configuration with symmetrical multiprocessing.

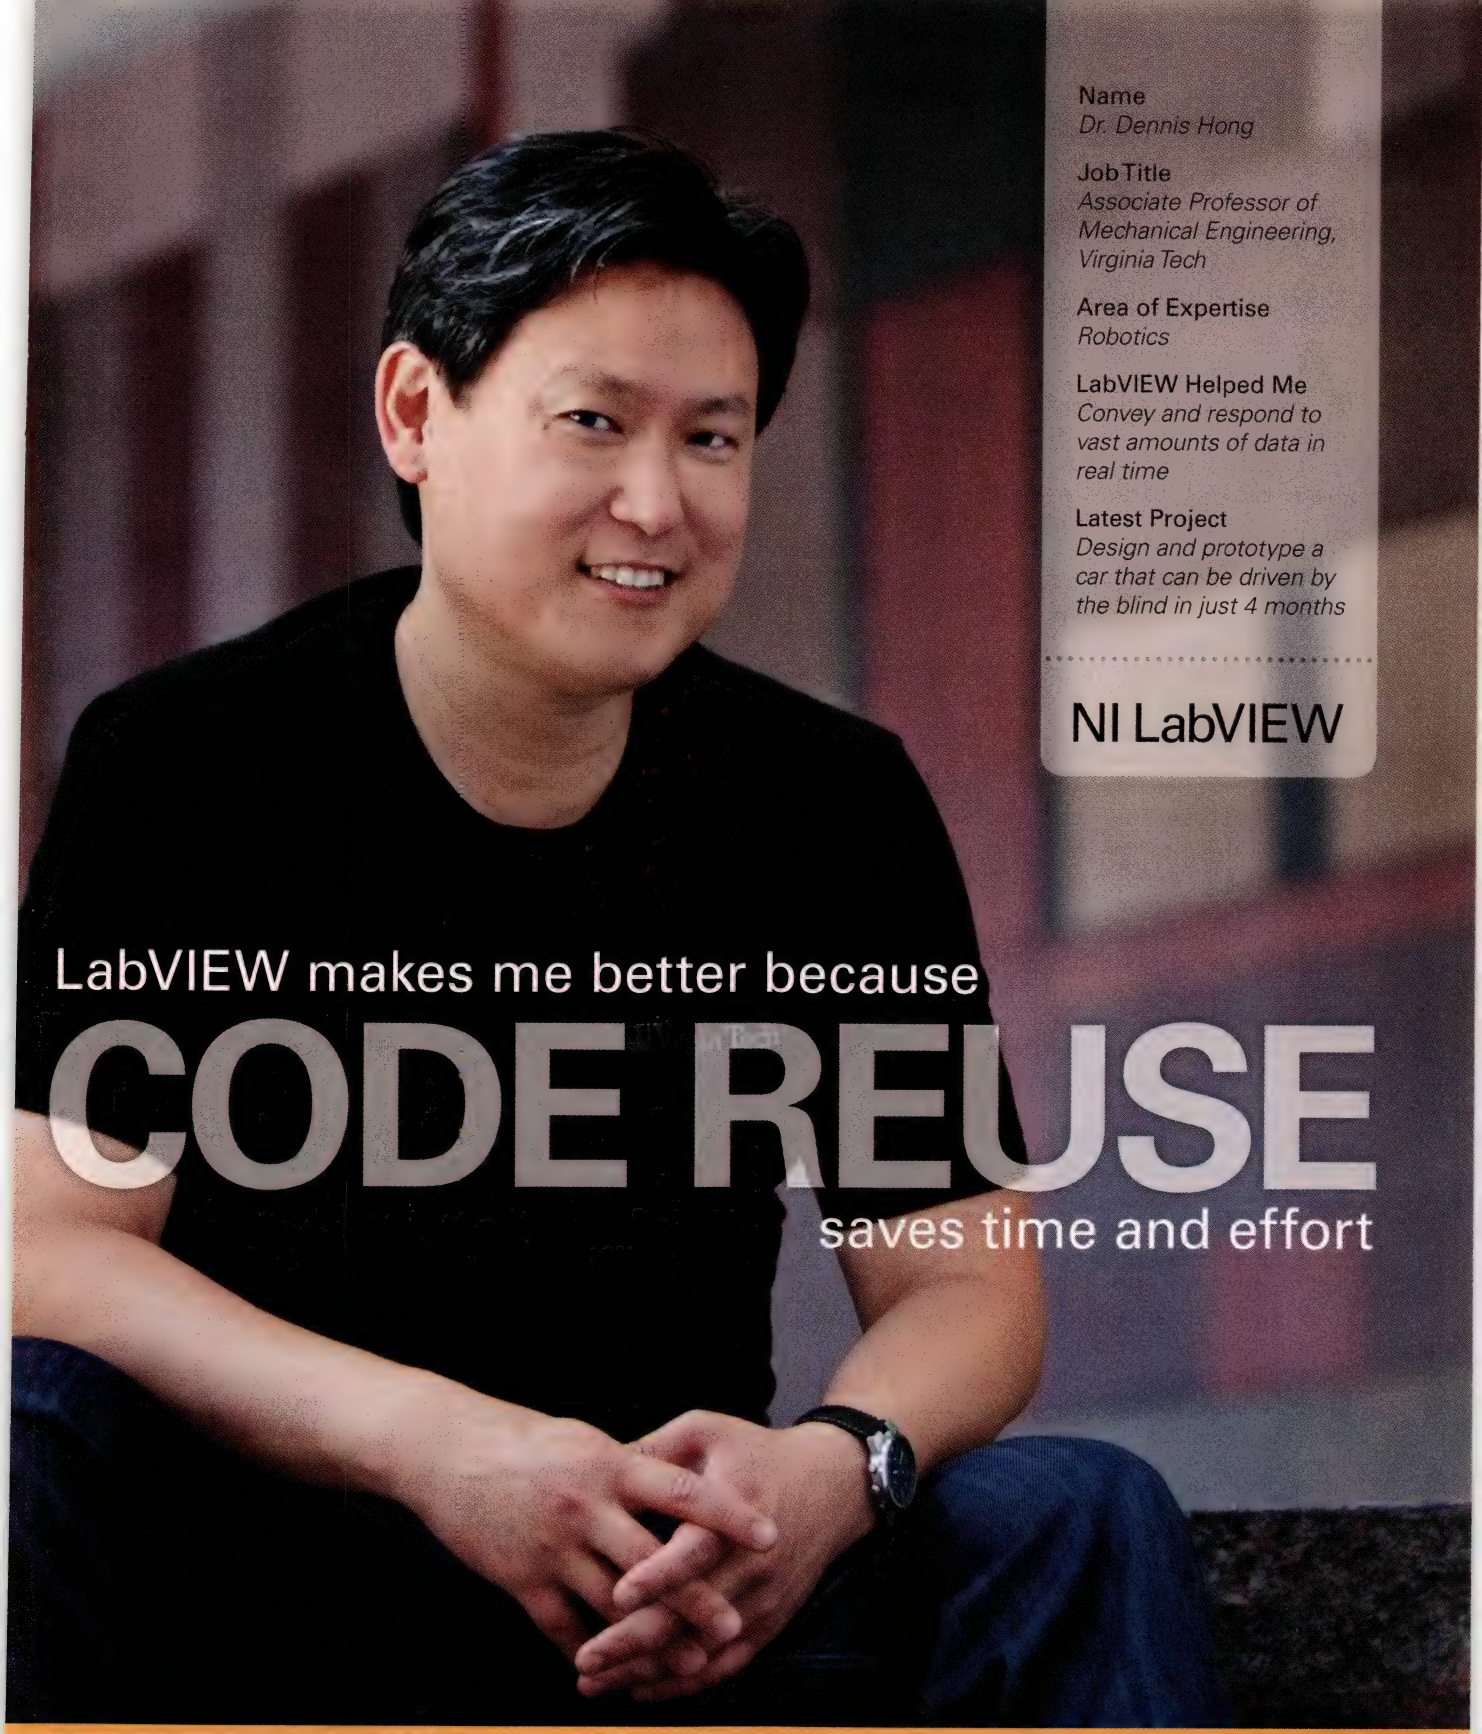

>> Find out how LabVIEW can make you better at **ni.com/labview/better** **800 453 6202**

NATIONAL INSTRUMENTS

System accurately characterizes, tests compliance of PCIe 3.0 receivers

Agilent Technologies has announced a test set for PCIe (Peripheral Component Interconnect Express) 3.0-receiver characterization. The Agilent PCIe 3.0 receiver-characterization test system provides receiver-tolerance test results and minimizes R&D effort. Using the test set, design and test engineers can accurately characterize and verify standards compliance of ASIC and chip-set receiver ports.

The PCI-SIG (Special Interest Group) recently released Revision 3.0 of the PCIe base specification. Many changes in the physical-layer spec relate to ensuring proper bit transmission at 8 billion transfers/sec over inexpensive PCBs (printed-circuit boards). The SIG uses receiver test to establish standards and requires them for demonstrating compliance with those standards.

Devices operating at these speeds present new measurement challenges, including optimizing transmitter and receiver equalization, calibration-channel design, and device-link training with 128/130-bit-coded pattern sequences. One of the most challenging requirements is implementing a new procedure for calibrating PCIe 3.0 stress conditions as a reference receiver would see them after applied equalization. Because the standards specify signals that exist only at inaccessible points, you must reconstruct the signals of interest using postprocessing signals that you can access.

Agilent based its approach on the J-BERT (jitter/bit-error-ratio tester) N4903B high-performance serial BERT, the N4916B de-emphasis signal converter, the 81150A pulse/function/arbitrary-noise generator, an Infiniium 90000 X series high-performance oscilloscope, new accessories, and new N5990A test-automation software.

The system provides stress-calibration software for accurate and repeatable receiver-test results, adjustable de-emphasis of precursor and postcursor conditions that precede and follow the system state you are examining, and the J-BERT's built-in PCIe 3.0-compliant jitter and sinusoidal-interference sources. It also offers new sweep capability for periodic jitter, new PCIe 3.0-compliant calibration channels, and higher efficiency, thanks to the new PCIe 3.0 stress-calibration software and link-training suite, which controls the J-BERT's pattern sequencer to bring the device under test into loop-back mode. It also features instrument protection resulting from the use of Agilent instruments that you can easily repurpose for accurate characterization of other multigigabit-per-second products, such as devices that support USB (Universal Serial bus) 3, SATA (serial advanced-technology attachment) 3, and QPI (QuickPath Interconnect) buses.

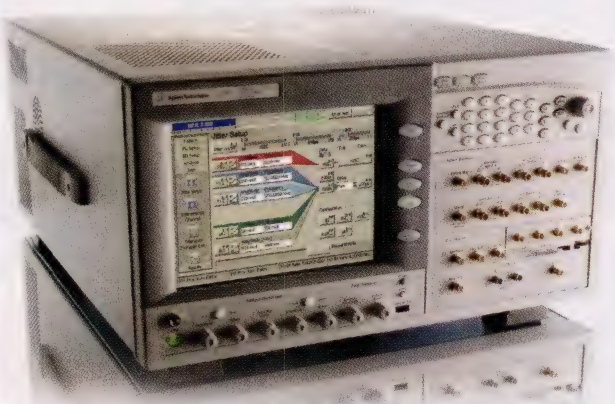

The N403B high-performance serial jitter/bit-error-ratio tester is at the heart of the manufacturer's new test set for characterizing and analyzing conformance of hardware for implementing the new 8 billion-transfer/sec PCIe 3.0 serial bus, which owes part of its speed to use of 128/130-bit error-correction coding instead of 8/10-bit coding.

The N4903B has integrated and calibrated jitter sources, including random and periodic jitter and differential-mode sinusoidal interference. The N4916B signal generator provides accurate, adjustable precursor- and postcursor-condition generation and an 81150A or 81160A pulse/function/arbitrary-noise generator for injecting common-mode sinusoidal interference. The system also includes N4915A-014 PCIe 3.0 calibration channels for compliant and reproducible channel conditions, N5990A-101 PCIe receiver-test software with new PCIe 3.0 stress calibration, the N5990A-301 PCIe 3.0 link-training wizard, and an Infiniium 90000 X series 16- to 32-GHz oscilloscope.

Prices for the 12.5-Gbps N4903B start at $192,000. The N4916B de-emphasis signal converter sells for $42,200. N4915A-014 PCIe 3.0 calibration channels cost $8150 each. The N5990A test-automation software Option 301 (link-training suite) and Option 101 (PCIe automated receiver test) sell for $9400 and $13,250, respectively. Users of N5990A-101 for PCIe 1.x and 2.x versions can use the $7200 Option N5990A-011 to upgrade to PCIe 3.0 capability. Prices for the 81150A pulse/function/arbitrary-noise generator start at $13,000 for a two-channel configuration.

—by Dan Strassberg
Agilent Technologies, www.agilent.com.

DILBERT By Scott Adams

>> Find out how LabVIEW can make you better at **ni.com/labview/better**

800 453 6202

NATIONAL INSTRUMENTS™

Midrange scopes feature pivoting displays

Two of the models in LeCroy's new six-member, midrange, four-channel, 400-MHz to 4-GHz WaveRunner 6 Zi oscilloscope family feature 40G samples/sec/channel in real-time acquisition when only one or two analog channels are in use. Each of the six 6 Zi units features a 12.1-in.-diagonal 800×1280-pixel-resolution color LCD that swivels into portrait or landscape orientation. The scopes' optional MSO (mixed-signal-oscilloscope) attachment adds as many as 36 logic-timing-analysis channels. The portrait mode provides more than 60% more room than landscape mode to vertically position or magnify the time-aligned traces. The pivoting screen can display 16 analog waveforms or as many as 15 analog and 36 MSO waveforms. The scopes' math capabilities can simultaneously generate multiple meaningful traces from each of the analog inputs. Derived traces include filtered waveforms and spectra obtained from input signals by use of Fourier transforms.

According to Dan Monopoli, product manager for LeCroy's midrange scopes, restricting the area that a waveform occupies on the screen of some competitive scopes forces the signal to use only a portion of its digitizer's dynamic range, noticeably degrading the waveform display. On the 6 Zi scopes, however, a waveform can span the digitizer's full dynamic range until you shrink the least-significant bit to less than the height of one display pixel.

Besides the six models that provide industry-standard 8-bit analog-waveform resolution, LeCroy is introducing two models with 12-bit resolution. These WaveRunner HROs (high-resolution-oscilloscopes) capture a maximum of 2G samples/sec/channel and provide bandwidths of 400 and 600 MHz. US suggested retail prices for 6 Zi units begin at less than $30,000.

—by Dan Strassberg

▷**LeCroy**, www.lecroy.com.

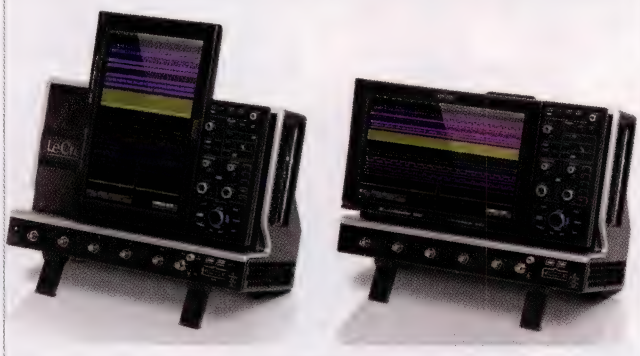

The pivoting screen immediately attracts attention to this midrange scope, but the instrument's unusual appearance is not its only noteworthy capability; a closer look at the specs reveals many significant features.

Thermal nanotape aims to cool hot spots

University-research consortium SRC (Semiconductor Research Corp) and researchers from Stanford University claim to have developed a combination of elements that yields a unique nanostructure material for packaging that they say offers longer device life for less cost than current solder methods. The partners note that current solder materials tend to degrade and fail due to heat and mechanical stress. Focusing on continued IC scaling, SRC and Stanford have researched materials that provide a thermal connectivity comparable to that of copper with the flexible compliance of foam. Their work has resulted in a nanostructured thermal tape, or nanotape, which conducts heat like a metal and allows the neighboring materials to expand and contract with temperature changes (references 1 and 2).

The work focuses on a combination of binder materials surrounding carbon nanotubes and a heat-sink contact. It looks at the benefits of using a nanostructured thermal-interface material, demonstrates carbon-nanotube growth on thermoelectric material, and measures the thermal-boundary resistances and heat capacity associated with carbon nanotube film.

Although the research targets ICs, other industries could gain from the thermal interface. For example, some automotive-related companies hope to recover electrical power from hot exhaust gases in cars and trucks using thermoelectric-energy converters, but reliable interfaces are a problem for this technology. Lead researcher Ken Goodson leads a major grant from the National Science Foundation (www.nsf.gov)/Department of Energy (www.doe.gov) Partnership on Thermoelectric Devices for Vehicle Applications, with the goal of transferring the SRC-funded interface work to vehicles.

Patents for the technology are pending. The next step in the research is to license the new methods and materials to advanced thermal-interface companies for perfection of the application. According to SRC, end users should benefit from the technology by 2014.

—by Suzanne Deffree

▶**Semiconductor Research Corp**, www.src.org.

▶**Stanford University**, www.stanford.edu.

03.17.11

REFERENCES

1. Gao, Y; AM Marconnet; MA Panzer; S Leblanc; S Dogbe; Y Ezzahri; A Shakouri; and KE Goodson, "Nanostructured Interfaces for Thermoelectronics," *Journal of Electronic Materials*, Volume 39, No. 9, 2010, http://bit.ly/i3pD6l.
2. Panzer, Matthew A; Hai M Duong; Jun Okawa; Junichiro Shiomi; Brian L Wardle; Shigeo Maruyama; and Kenneth E Goodson, "Temperature-Dependent Phonon Conduction and Nanotube Engagement in Metalized Single Wall Carbon Nanotube Films," *Nano Letters*, May 26, 2010, pg 2395, http://bit.ly/hgabrS.

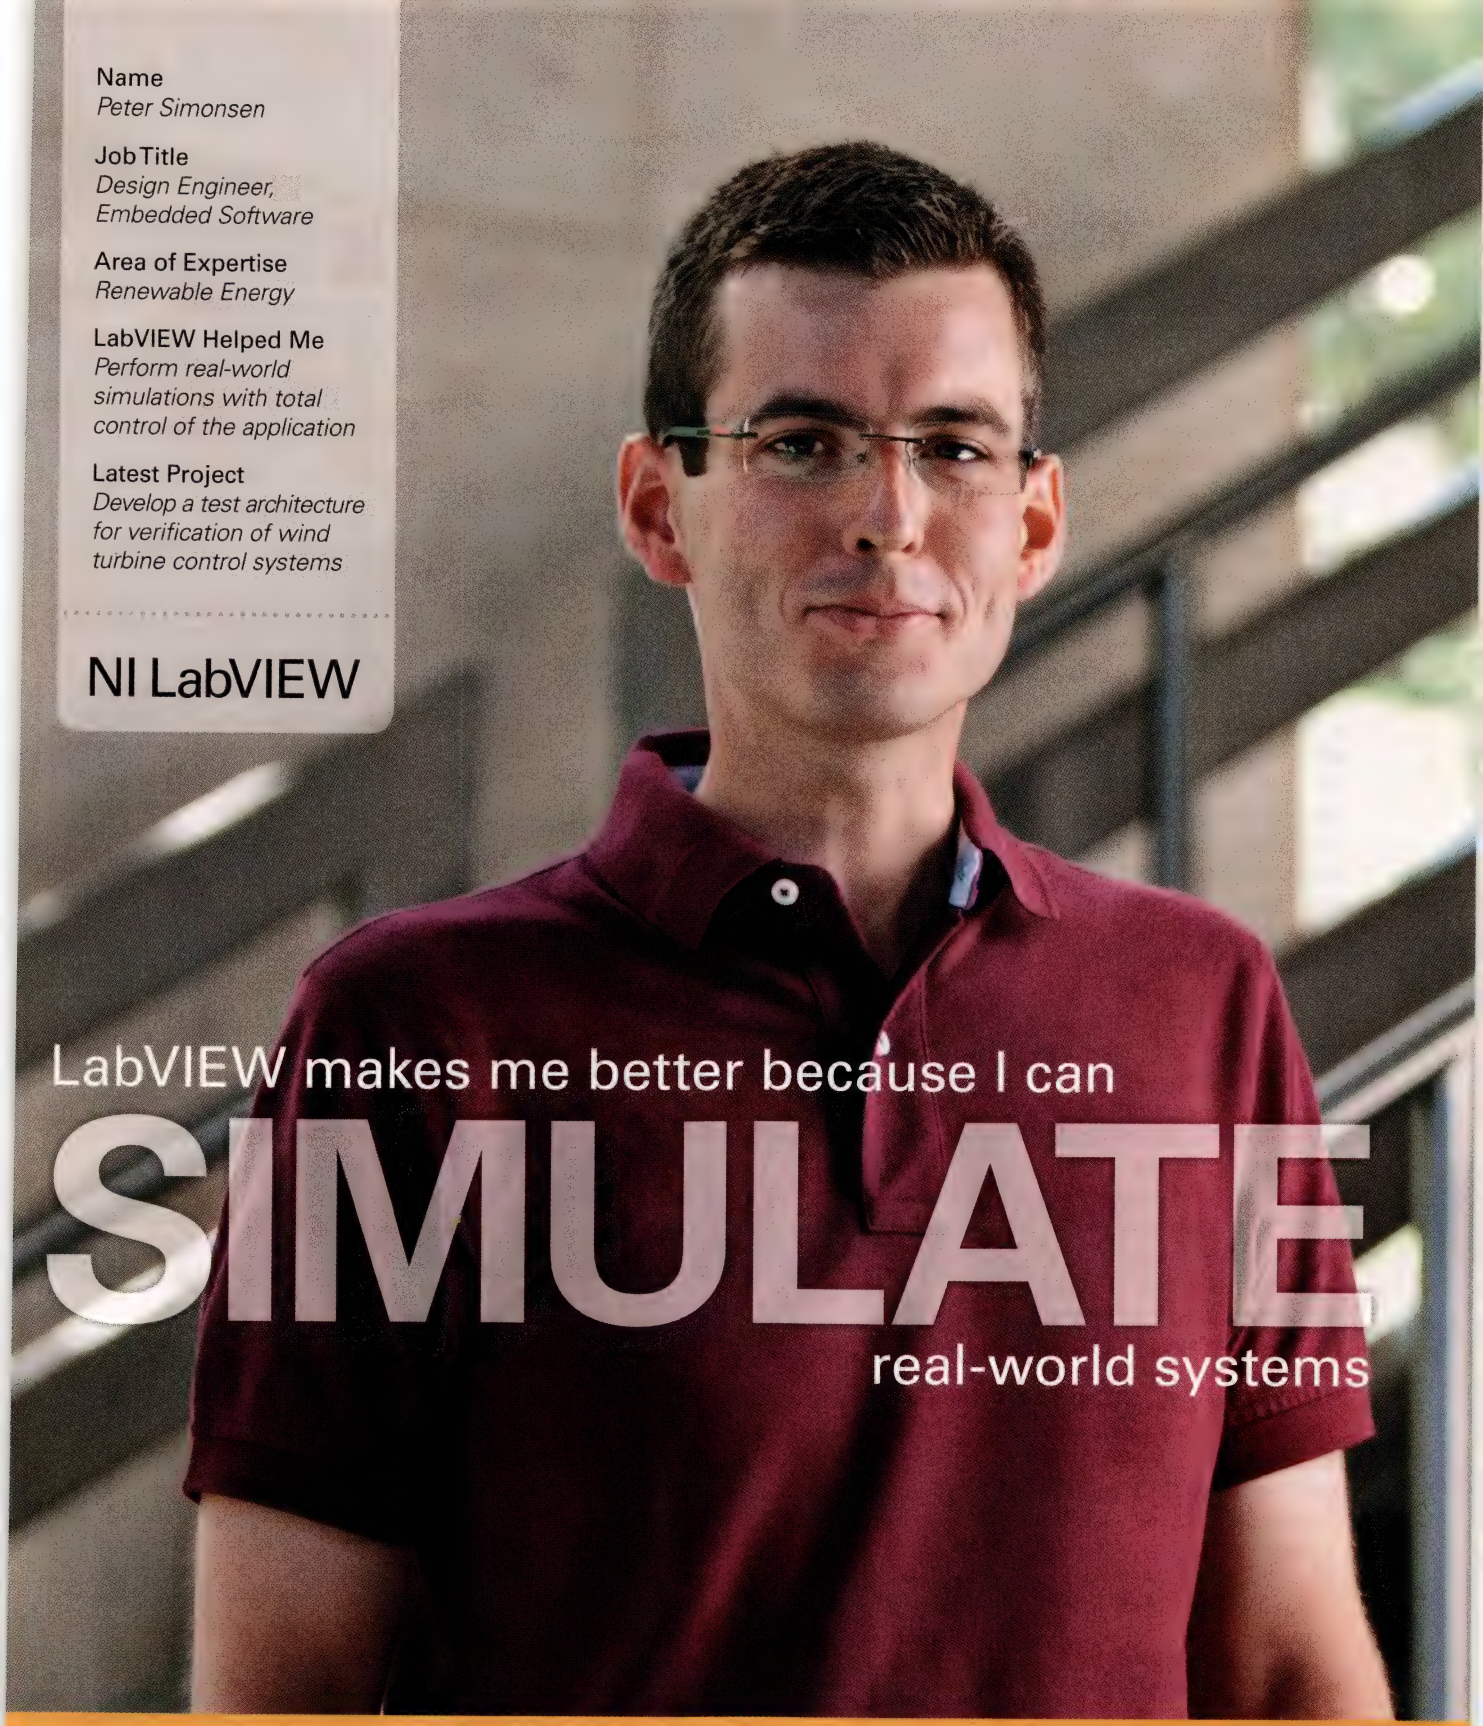

>> Find out how LabVIEW can make you better at **ni.com/labview/better** **800 453 6202**

NATIONAL INSTRUMENTS

Dual ARM Cortex-A9 MPCore features 28-nm, low-power programmable logic for high-end embedded systems

The new Xilinx Zynq-7000 family of EPP (extensible-processing-platform) devices targets high-end embedded-system applications, such as video surveillance, automotive driver assistance, next-generation wireless, and factory automation. The portfolio comprises four devices, each integrating a complete ARM Cortex-A9 MPCore-processor-based system with 28,000 to 235,000 cells of 28-nm, low-power programmable logic, the equivalent of 430,000 to 3.5 million ASIC gates.

The Zynq architecture differs from previous marriages of programmable logic and embedded processors by moving from an FPGA-centric platform to a processor-centric model. In a typical FPGA-centric development flow, engineers create their designs in HDLs (hardware-description languages), such as Verilog and VHDL (VHSIC HDL). Zynq supports a software-centric flow, in which developers work almost exclusively in high-level languages, such as C and C++. For software developers, the Zynq-7000 appears the same as a standard, fully featured ARM processor-based SOC (system on chip) that boots immediately at power-up and can run a variety of operating systems independently of the programmable logic.

Each Zynq-7000 EPP device embeds a dual 800-MHz ARM Cortex-A9 MPCore with a dedicated Neon coprocessor for media and signal processing that adds instructions for audio, video, 3-D graphics, and image and speech processing, along with a double-precision FPU (floating-point unit). The hard-wired processing system includes L1 and L2 caches, memory controllers, and commonly used peripherals. The devices integrate dual 12-bit ADCs that support sampling rates as high as 1M samples/sec on as many as 17 external-input analog channels.

The two largest devices in the Zynq family, the Zynq-7030 and Zynq-7040, have built-in multi-gigabit transceivers that operate as fast as 10.3125 Gbps and dedicated DSP resources that deliver 480 and 912 GMACS (billion multiply/accumulate operations per second) of peak performance, respectively. The two smaller devices, the Zynq-7010 and Zynq-7020, provide as much as 58 and 158 GMACS of peak DSP performance, respectively.

The EPP enables system architects, logic designers, and software developers to work in parallel within their familiar programming environments. System architects and software developers can initially define the system from the software perspective before determining what functions need offloading or acceleration in hardware. Software developers can leverage the Eclipse environment; Xilinx Platform Studio SDK (software-development kit); ARM Development Studio 5; ARM RealView Development Suite; and compilers, debuggers, and applications from vendors in the ARM Connected Community and Xilinx Alliance Program. These vendors include Lauterbach, Wind River, PetaLogix, The MathWorks, Mentor Graphics, Micrium, and MontaVista (www.lauterbach.com, www.windriver.com, www.petalogix.com, www.mathworks.com, www.mentor.com, www.micrium.com, and www.montavista.com).

A tight coupling of the processor system and programmable logic enables hardware/software co-design through AMBA-AXI (Advanced Microcontroller Bus Architecture-Advanced Extensible Interface) interconnects. The AMBA interfaces allow software developers to extract programmable logic as memory-mapped calls and to make more than 3000 internal connections to the programmable fabric. Configuration control allows for the creation of adaptive systems, in which developers can reprogram parts of the programmable logic through software to meet varying system operating environments. Hardware engineers can design the Zynq-7000 programmable fabric by using Xilinx's ISE (integrated-software-environment) design suite, which includes development tools and AMBA4 AXI4 plug-and-play IP (intellectual property) and BFMs (bus-functional models) to accelerate design and verification.

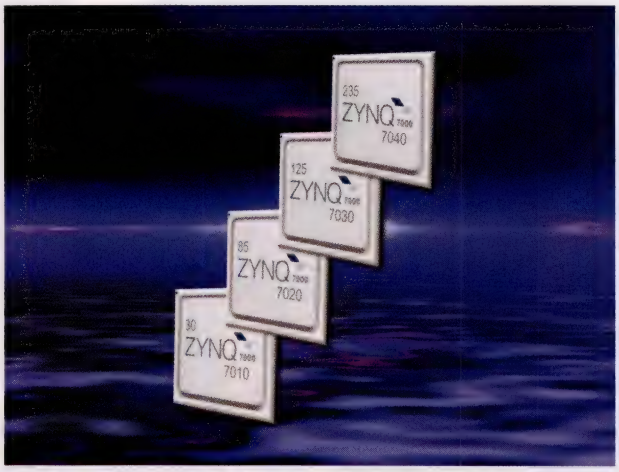

The Zynq-7000 family includes four devices that integrate a complete ARM Cortex-A9 MPCore-processor-based system with 28,000 to 235,000 cells of 28-nm low-power programmable logic.

> Developers can reprogram parts of the programmable logic through software for varying operating environments.

Customers can start evaluating the Zynq-7000 family by joining the Xilinx Early Access program, which Xilinx has limited to 100 participants. The company plans to release its first silicon devices for the second half of 2011, and engineering samples should become available in the first half of 2012. Prices vary, depending upon volume and device. Prices for the 7000 family will start at less than $15 (high volumes). A design kit will be available for $495. **—by Mike Demler**

Xilinx, www.xilinx.com.

MachXO2™ PLD

MachXO2.
DO-IT-ALL DEVICES DEMAND A DO-IT-ALL PLD.

With the MachXO2 PLD family, ycu now have everything you need to meet your low power, small form factor, high functionality and low price requirements for your next consumer or system design.

Access free reference designs and software, development kits and supporting literature at:
latticesemi.com/machxo2

Applications

- Smart Phone
- Digital Camera
- GPS
- Server
- Router
- Base Station

And more...

Benefits

- Instant-on, non-volatile
- Flexible logic from 256 LUTs to 7,000 LUTs
- Embedded memory up to 240 Kbits
- Ultra low power as low as 19µW
- On-chip user flash memory & I^2C, SPI, timer/counter
- Prices starting as low as $0.75

And more...

Lattice® Semiconductor Corporation

latticesemi.com/machxo2

FPGA-reference-design kit eases HDR-video-surveillance-camera development

Lattice Semiconductor's new HDR (high-dynamic-range)-60, a high-definition-videocamera-development system, employs the LatticeECP3 FPGA family. Lattice has partnered with Helion GmbH (www.helionvision.com), a spin-off of the Fraunhofer Institute in Germany, to integrate a preloaded plug-and-play evaluation ISP (image-signal-processing) pipeline into the HDR-60. The installed demo uses a subset of Helion's Ionos, a complete HDR ISP system from sensor to HDMI (high-definition-multimedia-interface)/DVI (digital-video-interface) display. The ISP delivers 1080p performance at 60 frames/sec with 2-D noise reduction and HDR of more than 120 dB per scene.

The HDR-60 kit includes a main PCB (printed-circuit board) with a Lattice ECP3-70 FPGA and a smaller NanoVesta sensor headboard with a 720p HDR sensor from Aptina (www.aptina.com). The company plans a 1080p HDR sensor for availability in the second quarter of this year. The device's image-processing functions include fast auto-exposure and an efficient automatic-white-balance algorithm. All of these functions take place in streaming mode through the FPGA without an external frame buffer. Onboard DDR-2 memory enables applications such as 3-D noise reduction, image stitching from multiple sensors, image rotation, and dewarping.

Lattice built the HDR-60 with a form factor to fit into commercially available camera housings, and it simultaneously supports two sensors. You can program the device by connecting to one of a pair of onboard USB (Universal Serial Bus) ports. The HDR-60 also includes an RJ-45 Ethernet port and a Broadcom Broadreach PHY (physical)-layer interface, offering support for Ethernet over RG-6 coaxial cable for distances as long as 700m at 100 Mbps for customers incorporating compression encoders into their designs. The HDR-60 targets rapid evaluation and prototyping of high-definition HDR videocameras for security and surveillance, traffic control, videoconferencing, and automotive applications.

Free schematics and Gerber layout files are available. The $399 kit includes an HDMI cable, an HDMI-to-DVI adapter, two standard USB programming cables, a universal power supply, and a quick-start guide.—**by Mike Demler**

Lattice Semiconductor Corp, www.latticesemi.com.

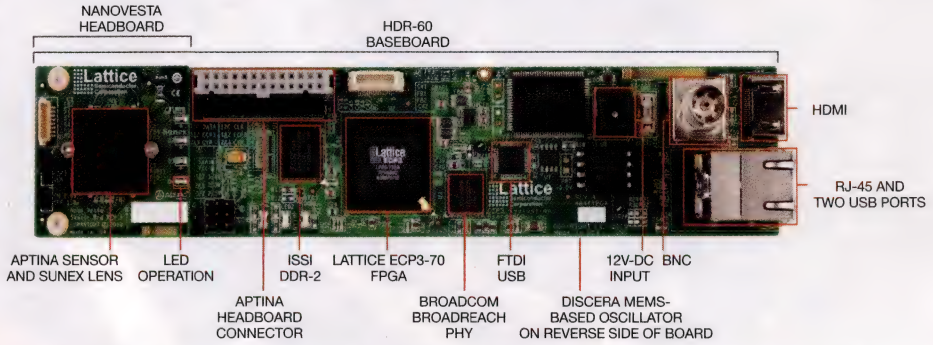

The Lattice HDR-60 reference-design kit includes a baseboard with an ECP3-70 FPGA and a NanoVesta headboard with an Aptina image sensor and Sunex lens.

Air hybrid cars could halve fuel consumption

Researchers in Sweden are exploring air, or pneumatic, hybrid automobiles as possible alternatives to electric hybrid cars, basing their findings on the fact that the action of braking releases energy. The researchers believe that manufacturers can save this energy for later use in the form of compressed air. The engine then gets extra power when the driver starts the car, which saves fuel by avoiding idling when at a standstill. The concept is similar to that in electric cars and electric hybrid cars, which use brake energy to power a generator that charges the batteries. According to Per Tunestål, a researcher in Combustion Engines at Lund University in Sweden, air hybrids would be less expensive to manufacture than electric and electric hybrid cars. Although these air hybrid vehicles are not yet in production, the step to commercialization need not be large, the researchers claim. The technology is attractive for stop-and-go and slow driving—buses in urban traffic, for example.

"My simulations show that buses in cities could reduce their fuel consumption by 60%," says Sasa Trajkovic, a doctoral student in Combustion Engines at Lund University. Trajkovic adds that cars could later reuse 48% of the brake energy, which the technology compresses and saves in a small air tank that connects to the engine. Thus, the degree of reuse for air hybrids could match that of today's electric hybrids.

> The technology is attractive for stop-and-go and slow driving.

The researchers developed a pneumatic hybrid-vehicle model in The MathWorks' (www.mathworks.com) Matlab and Simulink tools. The engine part of the vehicle model comprises engine data from Gamma Technologies' (www.gtisoft.com) industry-standard GT-Power model for engine simulation. Vehicle-drive-cycle simulations show that the technology could reduce the fuel consumption by as much as 58% below that of a conventional bus.

The researchers say that these simulated air hybrid engines require no expensive materials to manufacture and that they have smaller footprints than electric hybrids. The method works with petroleum, natural-gas, and diesel fuel. For more, go to http://bit.ly/hNeUyc.—**by Suzanne Deffree**

Lund University, www.lunduniversity.lu.se.

Windows.® Life without Walls.™ Dell recommends Windows 7.

Let your vision take shape.

DELL

The power to do more

Image provided by Autodesk and created with Autodesk® Revit software for building information modeling (BIM).

Unleash your creativity and imagine the possibilities. Then watch as they become reality with the power of Dell Precision™ workstations and Autodesk® BIM software.

Dell Precision™ workstations deliver the performance and graphics needed to run demanding applications with ease. Now your team can use data-rich modeling to evaluate new design options, predict building performance and communicate more productively.

- Blast through your workload faster than ever with the server-grade dual processor performance of a system powered by the Intel® Xeon® Processor 5600 Series. It's not just a workstation. It's an expert workbench.
- Genuine Windows® 7 Professional
- Scalable options – select systems are available with up to 192GB of memory* and 7.5TB of internal storage
- A full range of desktop, rack and mobile workstations to fit your needs
- ISV-certified for 95 leading applications, including Autodesk®
- Stay up and running with Next Business Day on-site warranty options or optional 24/7 Dell ProSupport™ covering applications from 25 ISV vendors, including Autodesk.

Get Equipped ▸

Find your ideal configuration online.
Go to Dell.com/smb/Vision or
call your sales rep at 1-800-624-6136.

intel Xeon inside™

Powerful. Intelligent.

*Actual capacity varies with preloaded material and operating environment and may be less. Significant system memory may be used to support graphics, depending on system memory size and other factors. Intel, the Intel logo, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries. Attn Warranties, One Dell Way, Round Rock, TX 78682 or visit dell.com/warranty. Remote Diagnosis is determination by online/phone technician of cause of issue, may involve customer access to inside of unit and multiple or extended sessions. If issue is covered by Limited Hardware Warranty and not resolved remotely, technician and/or part will be dispatched usually within 1 business day following completion of Remote Diagnosis. Availability varies. Other conditions apply.

VOICES

NXP's Rick Clemmer: short-range technology, long-range opportunity

Rick Clemmer, executive director, president, and chief executive officer of NXP Semiconductors (www.nxp.com), discusses the company after it released its results for the fourth quarter of 2010. With NXP's IPO (initial public offering) securely behind it and the soft landing of 2010 giving way to 2011, Clemmer discusses the improving investor view on the semiconductor industry, capacity investment, and the long-range opportunities that short-range technologies such as NFC (near-field communication) can provide. He also discusses NXP's strategies for thwarting counterfeiters.

When we last spoke, following NXP's third-quarter announcement, you noted that financial investors remained concerned about the overall health of the semiconductor industry. Are they still concerned?

A Most of that concern has gone away. If you look at the industry indices from October versus now, if you look at stock or the SMH [Semiconductor HOLDRs] indices, they are up. So I think that most of the concerns from general investors about the fundamentals of the semiconductor industry are now significantly behind us and they are focused on the differentiated players who are going to make a difference in winning in the market. It's totally a different environment from before.

Capacity in NXP wafer fabs averaged 97% in the fourth quarter of 2010 compared with 76% in the fourth quarter of 2009 and 99% in the third quarter of 2010. What capacity-expansion plans does NXP have in the works, and how will that expansion coexist with the company's Redesign Program, which, among other things, aims to streamline manufacturing?

A The most significant place that we need to add capacity is in our embedded-flash capabilities, where we need to be able to support our ID [identification] business, specifically our NFC business, as well as our 32-bit ARM-based microcontrollers. We've made capital investments in our SSMC [Systems on Silicon Manufacturing Co] in Singapore, a joint venture with TSMC [Taiwan Semiconductor Manufacturing Co], where we will double that capacity from the third quarter of 2010 to midyear 2011.

That [growth] gives us the ability to ensure that we can support customers on design wins and be able to drive that [growth]. In other areas, we are doing some balancing capacity, as well as some nominal investment in other manufacturing areas to be sure that we have the capacity to support our customers.

If you look at 2010, the $243 million or so we spent in capital expenditure was up three times on what we spent in 2009. We feel comfortable that we are making the investments to support our customers with the capacity.

Let's talk about NXP's efforts to thwart counterfeiting.

A We announced our authentication technology, which basically takes our ID technology to the next level. Looking at high-end goods, pharmaceuticals, and other things, [counterfeiting is] a big issue. KPMG [Klynveld Peat Marwick Goerdeler] recently did a study on electronic components that said that one of every 10 is counterfeit.

By having this authentication, it allows suppliers to ensure that it's their product. When it's a charger, you can confirm it with your authentication so someone can't use a knockoff. It protects your brand image and allows you to ensure that you have the sales associated with that image.

Obviously, there is some cost and implementation associated with that step, but we think the platform we are just beginning to roll out can be at a level at the same size of NFC two, three, or four years from now and represents a significant opportunity for us to continue to take the strength of our ID technology and continue to drive that technology on a broader basis.

A few weeks after Nokia announced plans for NFC, Apple is moving on NFC, a short-range-wireless technology, which developers often tout as a method for mobile payments. NXP has been a big supporter of NFC from day one. What does the Apple move mean for NFC, and how could it impact NXP?

A We can't talk about any individual company and what its plans are. I think it's safe to say all of the major handset and smartphone makers are working on it. When they deploy NFC and in which models is something we can't address. We can address it only after they have given us approval to talk about it, such as Google did on the Android platform and the fact that the Gingerbread platform has our NFC.

How do you feel, overall, about all of these companies' coming out in support of NFC? NXP, previously Philips Semiconductors, has been on NFC for a long time, but there haven't been many other major supporters.

A It's like in the movie, [*Field of Dreams*], "build it and they will come." Nobody came for about five years. Fortunately, the momentum is there now. Google's stepping up and trying to establish it on reference designs associated with the Android platform will allow the financial institutions and banks to really be able to drive that in conjunction with the carriers, and it will make a difference. As it begins wide deployment, its ease of use, with the "secure wallet," it will accelerate significantly.

—interview conducted and edited by Suzanne Deffree

RAQ's

Rarely Asked Questions

Strange stories from the call logs of Analog Devices

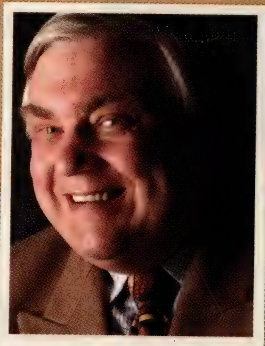

Contributing Writer James Bryant has been a European Applications Manager with Analog Devices since 1982. He holds a degree in Physics and Philosophy from the University of Leeds. He is also C.Eng., Eur. Eng., MIEE, and an FBIS. In addition to his passion for engineering, James is a radio ham and holds the call sign G4CLF.

Measuring the DC Parameters of Op Amps

Q: Modern op amps have very high open-loop gain and very low offsets. How can one possibly measure them? Or is that a closely held secret?

A: Remarkably easily, and absolutely not!

In the very first RAQ column, I praised the restaurant "La Cognette" in Issoudun in central France. I ate there again quite recently and, not long afterwards, at what is arguably the best restaurant in Africa, "Le Quartier Français" at Franschhoek in the South African winelands. At both places I had occasion to ask the chef how a dish was prepared, and at both I was willingly given all the details I could wish. The mediocre try to hide in secrecy, while the truly great wish to spread their knowledge.

Analog Devices is a world leader in analog technology. We believe that "If you cannot measure it, it's not science[1]" and are happy to share details of how we measure the parameters of our products.

Many textbooks include a diagram of an op amp with grounded inputs and a closed loop gain of 1–10 thousand driving a voltmeter that displays its offset amplified by 1–10 thousand. This technique works, provided care is taken to minimize thermoelectric voltages and the effects of bias currents, but it will only measure offset.

A slightly more complex arrangement, using a second op amp, allows the measurement of offset voltage, bias current, open-loop gain, common-mode rejection (CMR), and power supply rejection with minimal changes to the circuit and no low-level signal switching (which can introduce noise and errors). The addition of two more resistors and two more capacitors allows ac measurements as well. Furthermore, the second op amp need not be a high-performance device in order to make high-accuracy measurements (e.g., to measure very small offsets and very high gains).

The basic principle behind the circuit is that the auxiliary op amp provides feedback to the device under test (DUT), causing its output to be at a potential that forces the differential voltage at the auxiliary op amp's inputs to be (close to) zero.

While it is tempting to challenge readers to work out the fine details of the circuit's operation for all the different measurements, a detailed description is contained in "Simple Op Amp Measurements."

[1] In "The Door into Summer" by Robert A. Heinlein (Chapter 9), Dr. Twitchell says, "If you can't measure it, it's not science." This is an inversion of "If it's science, you can measure it." Popularly attributed to Lord Kelvin, this is a poor synopsis of what he actually said, which is nearer to Dr. Twitchell's observation. "In physical science, the first essential step in the direction of learning any subject is to find principles of numerical reckoning and practicable methods for measuring some quality connected with it. I often say that when you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind; it may be the beginning of knowledge, but you have scarcely in your thoughts advanced to the state of Science, whatever the matter may be." [Popular Lectures & Addresses, Vol. 1 "Electrical Units of Measurement" 1883-05-03]

To Learn More About Op Amp Measurements
http://dn.hotims.com/34927-101

Have a question involving a perplexing or unusual analog problem? Submit your question to: www.analog.com/askjames

For Analog Devices' Technical Support, Call 800-AnalogD

BY BONNIE BAKER

The electrical-overstress ghost

Customers sometimes tell me that my company has shipped a bad product. When I ask them what is wrong with the part, they say, "It just doesn't work." Many times, I find that customers use EOS (electrical-overstress) techniques to inadvertently cause this kind of problem.

One obvious source of EOS is an ESD (electrostatic-discharge) event, during which two bodies at different electrostatic potentials are in close proximity but hundreds or thousands of volts apart. Another EOS condition might occur if multiple power sources in a system turn on at different times. This situation may subject a pin or pins of a device in the system to an overvoltage condition. Still another overstress scenario is if a signal from the outside world or from another part of a system, under different power, appears on the input or the output of an amplifier.

Unknowingly, you may be relying on a device's internal ESD circuit to provide protection during an EOS event. You may find that an IC is perfectly operational before you apply power but that shortly after you apply the power and input signals, the IC becomes damaged. The EOS event may be so dramatic that the IC experiences extensive heating, causing both the die and the package material to melt (**Figure 1**).

IC amplifiers usually do not include protection against EOS. At best, the internal ESD-protection circuits may activate during an EOS event and provide sufficient protection. However, having on-chip ESD-protection circuitry does not protect it from all EOS situations.

A common EOS event can happen when applying the power-supply voltage to an amplifier after an input signal is present (**Figure 2**). If the 6.5-MHz rail-to-rail-I/O CMOS OPA374 op amp's input current is not limited, this power-up scenario can damage the amplifier's input ESD-protection circuitry. **Figure 2** shows results for this scenario simulated using TINA-TI SPICE software from Texas Instruments. In this case, the supply, V_{G2}, ramps from 0 to 5V in 50 msec; 5 msec after the supply begins its ramp, a 3.5V input signal, V_{G1}, appears at the amplifier's noninverting input. Shortly after 5 msec, the amplifier's input voltage is higher than the positive rail, turning on the noninverting input ESD diode. If the input source is low-impedance and can deliver current, a potentially harmful current can flow through the ESD diode. Using a series input resistor protects the input circuit from damage.

EOS events may last much longer than ESD events. When an EOS event occurs, the circuitry may conduct currents that lead to excessive heating and possible damage. If your system encounters EOS events, take measures to protect it with simple external protection circuitry. A minor design headache on the front end will save you from a major headache on the back end. EDN

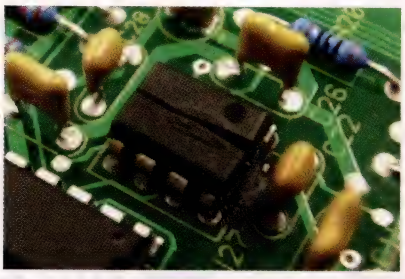

Figure 1 An EOS event can cause both the die and the package material to melt.

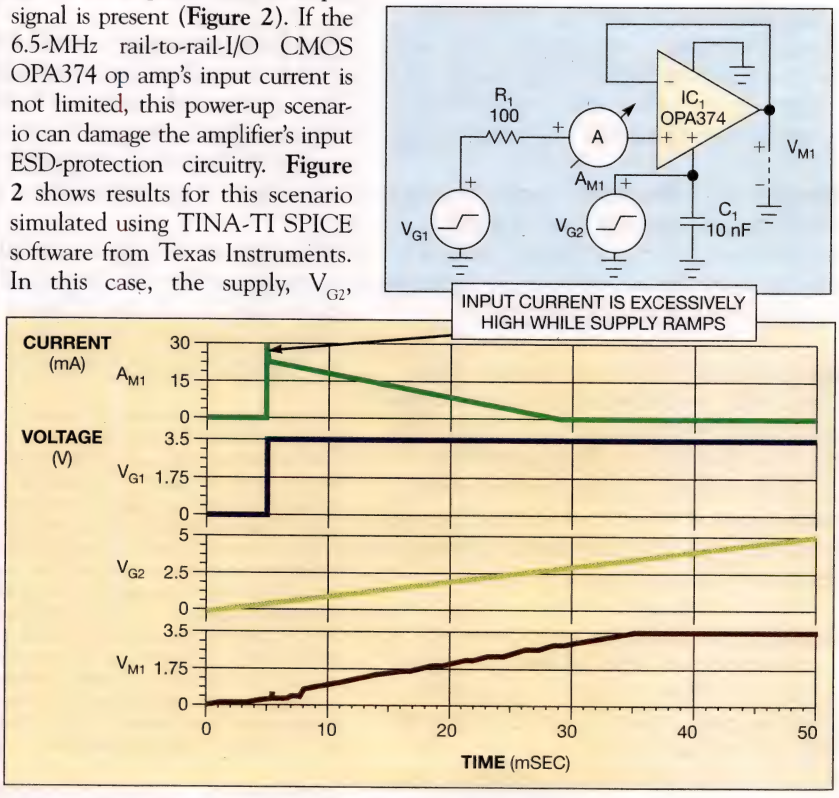

Figure 2 TINA-TI software from Texas Instruments simulated results for the scenario that can occur if the 6.5-MHz rail-to-rail-I/O CMOS OPA374 op amp's input current is not limited.

You Research. You Order. We Deliver.

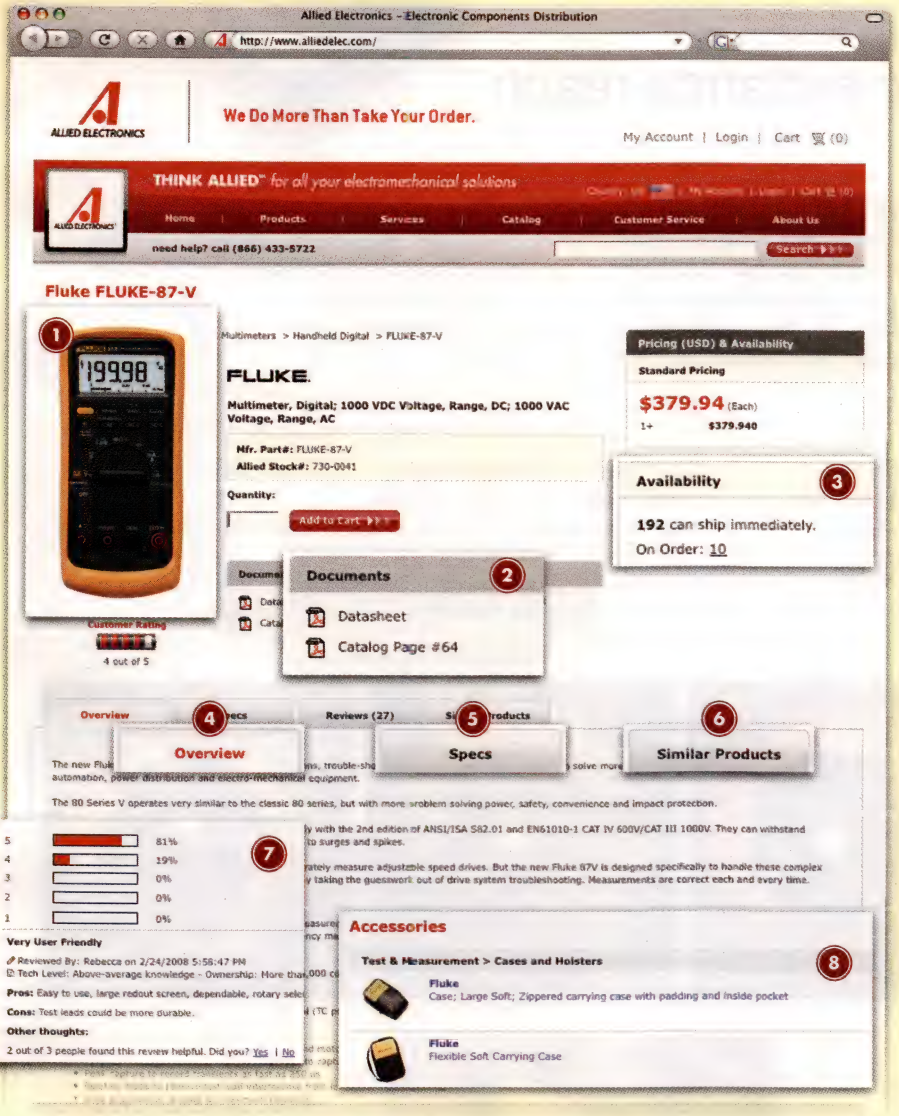

alliedelec.com

All the information you need at your fingertips!

Product Info

1. **Photos** - Detailed product photographs and expanded views help identify the exact part you need
2. **Datasheets** - Download manufacturer datasheets for in-depth product data
3. **Availability** - Check stock availability and stock on order, so you know exactly when your product will ship
4. **Overview** - Brief description of the product and its capabilities
5. **Specs** - Quick view of the product's technical specifications
6. **Similar Products** - Easily find similar products when the search result doesn't quite meet your needs
7. **Ratings** - View product reviews and ratings from other users
8. **Recommended Accessories** - Compatible accessories are available on thousands of products

Browse more than 2 million electrocomponent product solutions, and get all the information you need to place your next order.

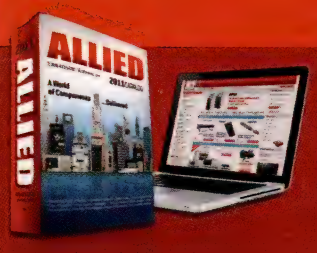

THINK ALLIED℠

alliedelec.com | 1.800.433.5700

ALLIED ELECTRONICS®

NEDA Supplier Authorized Distributor

PRY FURTHER AT EDN.COM

Go to www.edn.com/pryingeyes for past Prying Eyes write-ups.

Remote phosphor expands reach of LED light

The 100-year reign of the incandescent light bulb will end in 2012 as the Energy Independence and Security Act of 2007 mandates the phase-out of most incandescent bulbs in favor of more efficacious forms of lighting. New forms of lighting will also need to meet consumers' expectations for light color, dimming performance, lifetime, and the light pattern that a bulb casts.

To evaluate light patterns, I used a home-made jig that's basically a slanted rim—a conical section of light cardboard—that encircles a light bulb, reflecting its light over the 360° around the bulb.

Here's the familiar global-light distribution of an incandescent bulb.

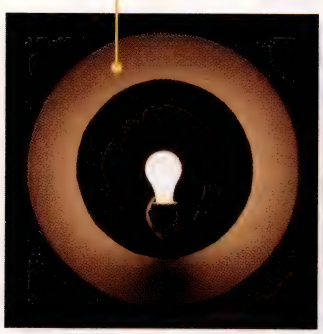

The CFL (compact fluorescent light) also has a fairly universal light distribution, although slightly less than that of the incandescent.

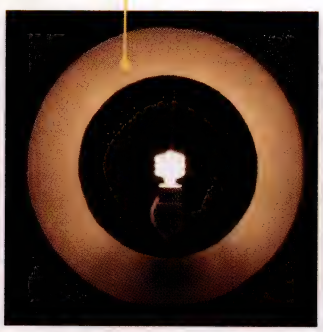

The snow-cone LED light was less than half of a semispherical pattern, which you'd expect: The LEDs rest on a flat surface, pointing up with little to no secondary optics to direct the light to the sides or down.

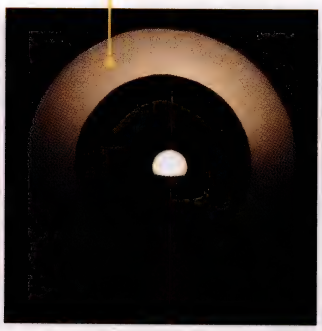

Now, take a look at the Philips light. It is virtually the same as the incandescent light.

How does the Philips bulb create such a spherical light pattern? Let's take a closer look at the bulb. Rather than mounting the LEDs horizontally on the base of the bulb, the LEDs mount vertically with their light shining out and striking the yellow case that surrounds each LED section.

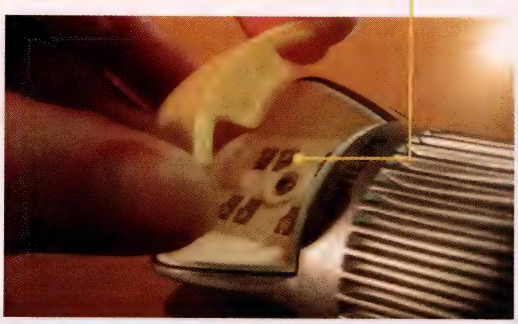

The yellow plastic is a remote phosphor. White LEDs comprise blue LEDs with a dollop of phosphor directly on top of and touching the blue-LED emitter. By using a remote, or secondary, phosphor, the bulb can use the remote phosphor's characteristic of unidirectional and uniform light emission, rather than as a point source. After carefully removing the yellow plastic remote phosphor and powering up the bulb, you can see the blue LEDs.

MathWorks®
Accelerating the pace of engineering and science

↘
Find it at
mathworks.com/accelerate
datasheet
video example
trial request

RUN MATLAB PROGRAMS IN PARALLEL

with
Parallel Computing Toolbox™

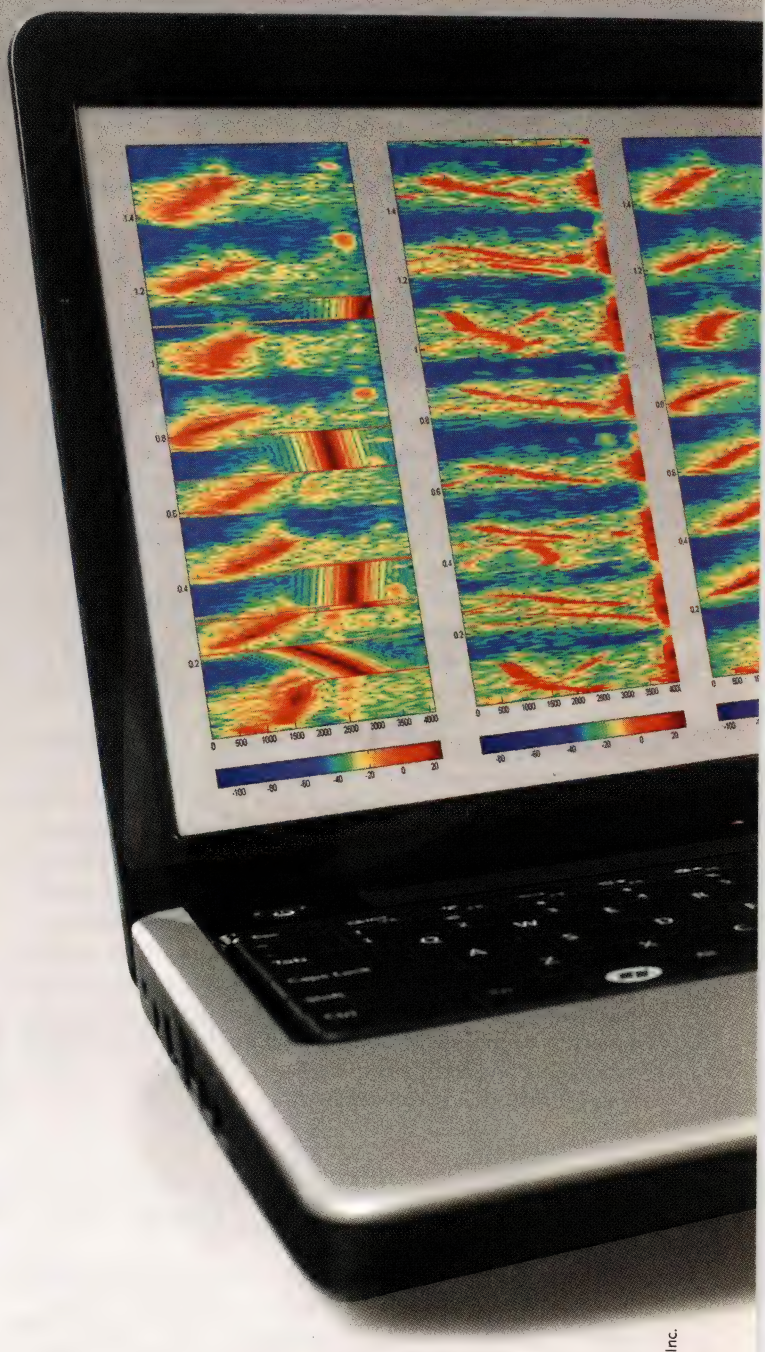

If you can write a FOR-loop, you can write a MATLAB parallel program for high-performance computing. Parallel Computing Toolbox lets you quickly adapt MATLAB programs to run on your multicore computer, GPU or cluster, with features for task-parallel and data-parallel programs.

MATLAB® & SIMULINK®

UBM Electronics
DesignNews EDN Test & Measurement World
MECHATRONICS IN DESIGN
FRESH IDEAS ON INTEGRATING MECHANICAL SYSTEMS, ELECTRONICS, CONTROL SYSTEMS, AND SOFTWARE IN DESIGN

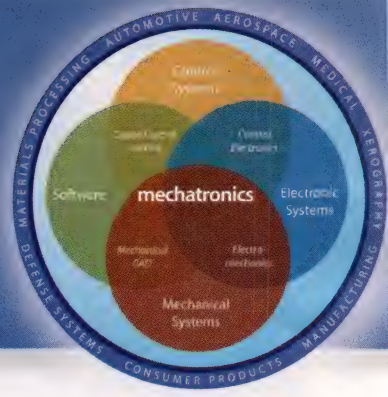

Impedance: an essential multidisciplinary concept

The impedance principle is core to connecting mechanical subsystems, but mechanical-engineering education and practice have neglected it.

By Kevin C Craig, PhD

Electrical engineers apply essential engineering concepts, such as impedance, all the time, whereas mechanical engineers most likely have not heard of it. No one, however, questions its utility in understanding and predicting how multidisciplinary systems perform when they interconnect. But what is impedance?

Imagine connecting two simple, mechanical, first-order subsystems. The first system might represent a machine-tool slide that a motor is positioning by providing force f_{i1}; the second system could be a velocity-measuring device you want to attach to the slide to measure its speed.

You can correctly predict the dynamic behavior of the interconnected subsystems in one of three ways. The most common way, although incorrect, is to separately derive each system-transfer function and then multiply them together, observing that $x_{o1}=x_{i2}$ when the systems connect. This approach completely ignores the fact that the second subsystem loads—that is, draws energy from—the first subsystem. Yet controls courses typically present the approach in this way, connecting transfer functions in series and multiplied together. They rarely mention loading. The first correct approach is to consider the system as a whole, with connected subsystems, and analyze it from scratch. The following **equation** provides the correct result:

$$\frac{x_{o2}}{f_{i1}}=\frac{\dfrac{b_2}{(b_1+b_2)k_2}}{\dfrac{m_1b_2}{(b_1+b_2)k_2}s^2+\dfrac{m_1k_2+b_1b_2}{(b_1+b_2)k_2}s+1}.$$

Each subsystem here has an input port and an output port, with two variables, force and velocity, at each port. The product of force and velocity is instantaneous power. The second approach to describe the dynamic behavior of this system is to derive a 2×2 transfer-function matrix showing the relationship among these four variables. The following **equation** shows the result of such an analysis:

$$\begin{pmatrix} f_{i1} \\ x_{i1} \end{pmatrix}=\begin{pmatrix} 1 & m_1s^2+b_1s \\ 0 & 1 \end{pmatrix}\begin{pmatrix} f_{o1} \\ x_{o1} \end{pmatrix} \text{ and } \begin{pmatrix} f_{i2} \\ x_{i2} \end{pmatrix}=\begin{pmatrix} 1 & k_2 \\ \dfrac{1}{b_2s} & \dfrac{b_2s+k_2}{b_2s} \end{pmatrix}\begin{pmatrix} f_{o2} \\ x_{o2} \end{pmatrix}.$$

Recognizing that $f_{o1}=f_{i2}$ and $x_{o1}=x_{i2}$ when the systems connect, you can multiply these two matrices together to predict the correct complete system behavior.

These two approaches work well if you have a mathematical model for each subsystem. Sometimes, however, you have only the actual hardware without any models available; you must still connect the subsystems for a well-performing, complete system.

How can an engineer ensure that this connection will happen without the benefit of modeling? The answer is the impedance concept; impedance is the ratio of the two variables whose product is power. When two systems connect, you must modify the upstream ideal, unloaded, subsystem-transfer function due to the loading effect that the downstream subsystem connection causes. That modification is to multiply the ideal upstream transfer function by the following **equation** to produce the loaded transfer function:

$$\left(\frac{1}{1+\dfrac{Z_{o1}(s)}{Z_{i2}(s)}}\right), \text{ where } Z_{o1}(s)=\left.\frac{x_{o1}}{f_{o1}}\right|_{f_{i1}=0} \text{ and } Z_{i2}(s)=\left.\frac{x_{i2}}{f_{i2}}\right|_{f_{o2}=0}.$$

Here, Z_{o1} is the output impedance of Subsystem 1 and Z_{i2} is the input impedance of Subsystem 2. If Z_{o1} is much less than Z_{i2} over the frequency range of interest, then loading effects are negligible. This approach gives the exact result, as do the previous two approaches. This approach is important because you can experimentally measure the subsystem's ideal upstream and downstream transfer functions, along with Z_{o1} and Z_{i2}. Once you measure them, you can use the results to properly predict the behavior of the assembled system. You can also use this approach with mathematical models to predict the correct behavior. EDN

Multi-lane Jitter Tolerance running too slow?

Let Centellax Save You Time

If testing jitter tolerance in multi-lane devices or high channel count SERDES with a single serial BERT gives your test technician plenty of time for more exciting things, like watching plants grow, look to Centellax to speed things up. **Our Parallel Channel BERT measures Jitter Tolerance on ALL lanes simultaneously**, at a surprisingly low cost that won't break your capital budget.

Speed Up Your Testing

Minimize cable loss with remote pattern generator and error detector heads that can be mounted close to your device under test. 4-tap de-emphasis is integrated directly in the pattern generator head, eliminating the need (and cost) of a dedicated de-emphasis processor. Our Signal Integrity Studio software ties it all together, including a parallel jitter tolerance test with unlimited channel count.

When you are ready to speed up your testing, come see us at www.centellax.com.

CENTELLAX
SIGNAL INTEGRITY MEASUREMENT

Affordability without compromise

For more information on Centellax products visit **www.centellax.com**

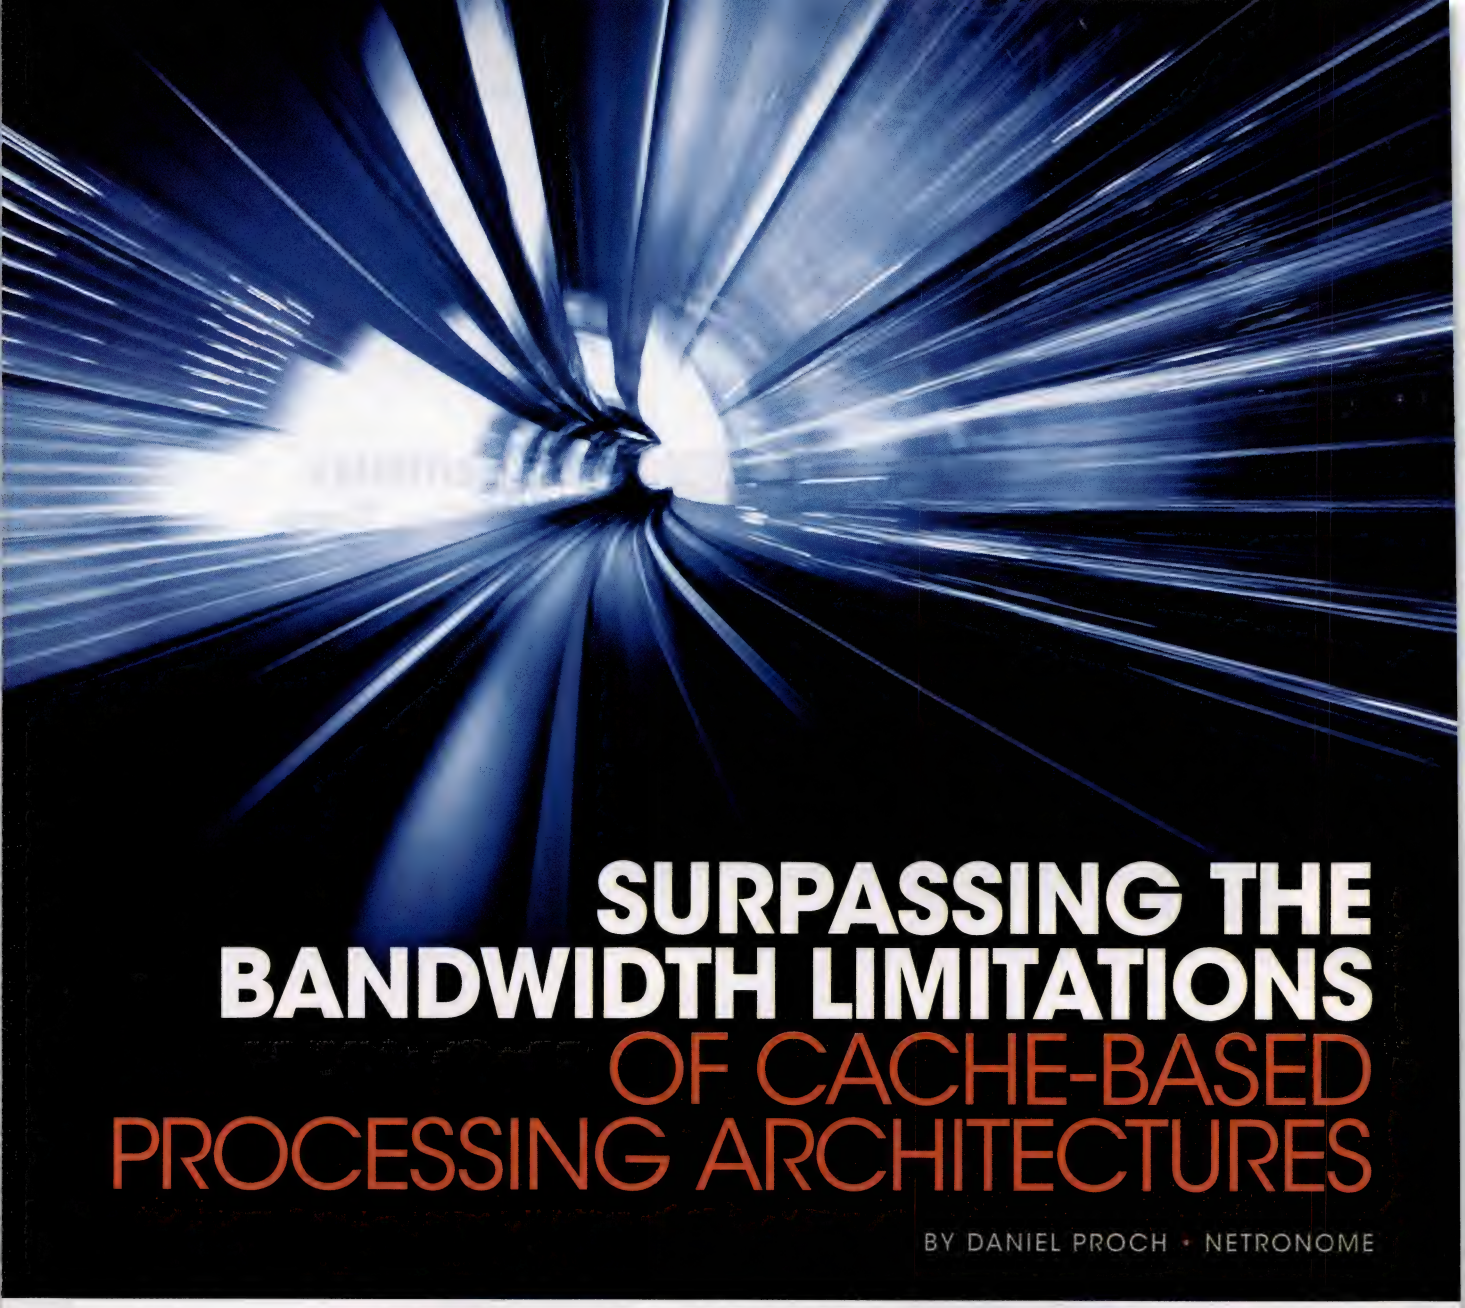

SURPASSING THE BANDWIDTH LIMITATIONS OF CACHE-BASED PROCESSING ARCHITECTURES

BY DANIEL PROCH • NETRONOME

AS BANDWIDTH INCREASES AND SECURITY THREATS EVOLVE, NETWORK INFRASTRUCTURES MUST SUPPORT 10 AND 40 GBPS OF THROUGHPUT, WITH DEEPER INTELLIGENCE. TRADITIONAL CACHE-BASED PROCESSORS DON'T SUIT USE IN DATA-PLANE PROCESSING ONCE IT SCALES TO MILLIONS OF FLOWS PER SECOND. AN ALTERNATIVE TIGHTLY COUPLES NETWORK FLOW PROCESSORS AND x86 CPU CORES.

The amount of network traffic in today's wired and wireless infrastructures continues to rise at dramatic rates to keep up with the demand for IP (Internet Protocol)-based voice, video, and data services and applications (**Figure 1**). Cisco estimates that annual global IP traffic will increase fourfold by 2014, growing from 176 exabytes/year to three-quarters of a zettabyte—that is, 767 exabytes (**Reference 1**). The primary drivers for this growth are video services and mobile data.

Video, such as TV, video on demand, Internet video, and P2P (peer to peer), will exceed 91% of global consumer network traffic within four years. Internet-based video will grow from 33% to almost 60% of Internet-data traffic, in the process surmounting P2P as the primary video contributor and comprising the equivalent of 12 billion DVDs of data per year. Mobile-data traffic, although still a smaller individual category, will double every year, increasing

Hello future.

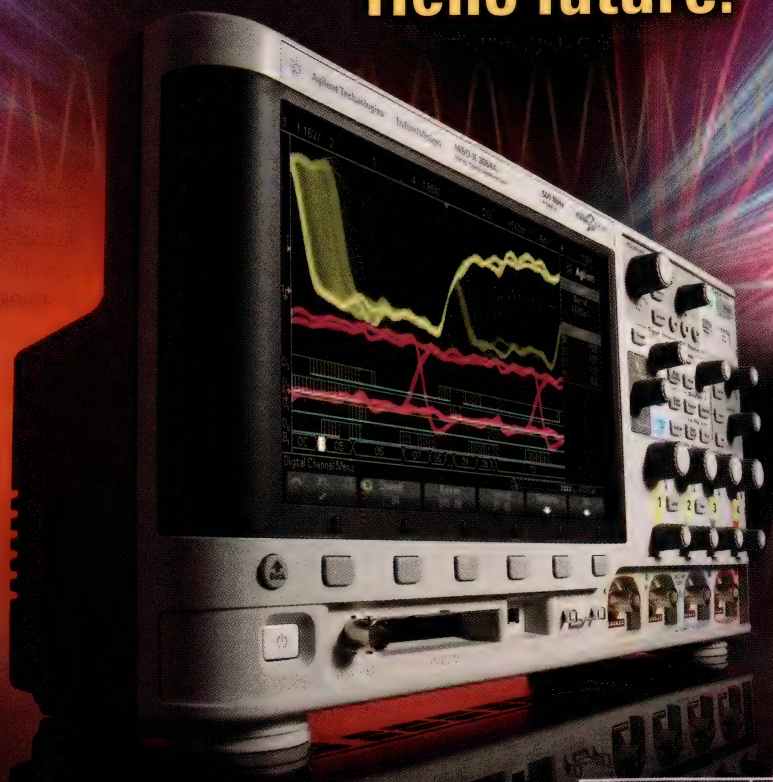

Oscilloscopes Redefined

Starting at
$1,230*

	Agilent 2000 X-Series (MSO and DSO)	Popular Competitor's Series (DSO)	Agilent 3000 X-Series (MSO and DSO)	Popular Competitor's Series (MSO and DSO)
Bandwidth (MHz)	70, 100, 200	50, 70, 100, 200	100, 200, 350, 500	100, 200
Max sample rate	2 GSa/s	2 GSa/s	4 GSa/s	1 GSa/s
Max memory depth	100 kpts	2.5 kpts	4 Mpts	1 Mpt
Max update rate (waveforms/sec)	50,000	200**	1,000,000	5,000
Fully upgradable	Yes	No	Yes	No
Function Generator	Yes	No	Yes	No
Notes:	**Refer to Agilent Pub 5989-7885EN for update rate measurements Data for competitive oscilloscopes from competitor publications 3GW-25645-0 and 3GW-22048-1			

Agilent and our Distributor Network
Right Instrument.
Right Expertise.
Delivered Right Now.

TestEQUITY
800-732-3457
www.testequity.com/agilent

See the difference today.
www.testequity.com/Agilent_Scope

*All prices are in USD and subject to change

Agilent Technologies

39 times in this same four-year period. And P2P traffic, although no longer comprising the most voluminous traffic type by 2014, will still be substantial as a percentage of overall network data.

An important implication of this rapid growth in network throughput exists from the perspective of designing and developing products for data-center and carrier environments. The GbE (gigabit-Ethernet) infrastructure that vendors commonly deploy will rapidly saturate, forcing network architects to quickly move to the 10-Gbps Ethernet successor. Vendors will also introduce follow-on 40- and 100-Gbps interfaces in the coming years, which will eventually become commonplace in switching, routing, and network-security products.

WHY FLOWS MATTER

More users and more applications are the fundamental driving forces of this dramatic increase in network throughput. This combination ultimately results in more individual "conversations," or "flows," traversing the network at any time. A flow is a unidirectional sequence of packets, all sharing a set of common packet-header values. Two packet-header fields, a source and destination IP-address combination, or as many as 11 Layer 2 through Layer 4 header values can identify a flow (**Figure 2**).

Most network equipment employing ASICs or fixed-function network processors, including Ethernet switches and IP routers, processes traffic based solely on the information contained in datagram headers. These devices process traffic packet by packet, keeping no in-memory information, or "state," of previous packets after each forwarding decision.

Network architects also deploy an array of security applications to protect their critical enterprise and carrier resources. Security-enhancing applications include virus scanning, firewalls, intrusion-detection and -prevention systems, DDOS (distributed-denial-of-service)-mitigation programs, DLP (data-loss-prevention) and test-and-measurement utilities, and network-forensics systems. These applications work almost entirely by implementing DPI (deep-packet inspection) and flow analysis, looking for known network patterns and, upon finding them, blocking or recording them. With the need for application awareness, security processing, and DPI, the amount of processing power for these computationally intensive applications grows exponentially with increases in line rates.

AT A GLANCE

- Cisco predicts a fourfold increase in global IP (Internet Protocol) traffic by 2014, with video and wireless services leading the charge.
- Increasingly elaborate network-security approaches require sophisticated packet processing, a high available-instruction-per-packet rate, and stateful flow management, all at 10 Gbps and higher speeds.
- General-purpose multicore CPUs are ineffective at data-plane processing for networking applications because the data in these applications is rarely spatially or temporally associative, and CPUs' caches are therefore too small to meaningfully encompass it.
- Network flow processors use multiple techniques to hide memory latencies, therefore providing more efficient memory-bandwidth usage than that of general-purpose processors.
- A heterogeneous multicore architecture tightly couples network-flow-processor cores with general-purpose multicore x86 systems over a high-speed virtualized PCIe (Peripheral Component Interconnect Express) datapath.

Maintaining the network state on all flows passing through a system is a critical requirement for all of these intelligent applications. Rather than implementing simple packet-based processing, security systems require sophisticated packet and security processing, along with a high available-instruction-per-packet rate and stateful management of flows at 10 Gbps and higher speeds.

EXAMPLE APPLICATIONS

Considering the evolution of today's threat landscape, numerous applications would prove ineffective without flow-based stateful processing of network traffic at the line rate. Cyber-security, lawful-interception, and traffic-management applications using DPI and behavioral-analysis techniques must retain a per-flow state because reliable analysis often requires seeing across individual packet boundaries to identify protocols and applications. These applications may also use heuristics or behavioral analysis to reliably detect applications or protocols even if advanced obfuscation or encryption techniques are in use.

As attacks become more sophisticated and attackers become better organized, intrusion-detection and -prevention systems rely on flow processing with many states. Modern attacks use invasion techniques, such as spreading malicious traffic across packet boundaries, payloads, and even IP fragments, to avoid detection. For example, Snort, a popular open-source intrusion-detection and -prevention application, includes a preprocessing module that reassembles an entire TCP (Transmission Control Protocol) flow to run signature-based rules against the entire connection payload, rather than simply examining traffic on a per-packet basis.

Network forensics, data-loss prevention, and antivirus applications, whether host- or network-based, terminate connections at the TCP layer, parse the application protocol, such as HTTP (Hypertext Transfer Protocol), SMTP (Simple Mail Transfer Protocol), P2P, and others, and even reassemble entire file attachments to scan for threats and monitor for confidentiality breaches.

The emergence of stateful next-generation firewalls, devices that integrate traditional firewall and network-intrusion-prevention capabilities, has recently caused a major stir in the market. The essential requirements for an effective next-generation firewall include the ability to identify applications regardless of port, protocol, or encryption scheme; to provide visibility and control over applications; and to accurately identify users to provide real-time protection against a variety of threats, including those at the application layer. A next-generation firewall retains significant attributes of each connection in memory, in which application identification and security processing happen at the

beginning of the flow. The firewalls then use the flow state to process the session as a means of increasing performance.

FLOW CHALLENGES

As networks' traffic and bandwidth increase, building these networks becomes an increasingly memory-intensive challenge. Processing huge volumes of traffic at high instruction rates and maintaining an accurate tracking of flows require large amounts of memory for a state to remain across all of the packets in the flow.

Analysis of packet captures of real-world network-backbone links enables further investigation of flow-based forwarding challenges—specifically, the relationship between network throughput, packet size, and flow length in an effort to understand the mean time between packets in a flow (**Figure 3**). From such information, architects can derive the system memory for stateful flow processing at 10 Gbps and beyond.

THE STATE REQUIRED TO PROCESS FLOWS INCREASES LINEARLY WITH AN INCREASE IN TRAFFIC IN NETWORKS WITH SIMILAR TRAFFIC PROFILES.

It can be shown that the state required to process flows increases linearly with an increase in traffic in networks with similar traffic profiles. Analysis also reveals that the interpacket time within a flow is almost entirely due to application delay and tributary network speed. Transactional and signaling flows tend to be shorter and have greater application delays than do flows from streaming applications (**Reference 2**). The phenomenon is not due to network-backbone links. Factors such as network transmission and delay are also statistically irrelevant to interpacket times.

You might expect that packets that are sent back to back would be a few microseconds apart, but the data shows that flows are highly temporally dissociative. Average flow interpacket times can be a second or longer. Analysis also shows that bandwidth has no significant effect on flow interpacket time. Bandwidth does dramatically affect other aspects related to flow processing, however. As throughputs increase, the total number of flows increases linearly with throughput, as

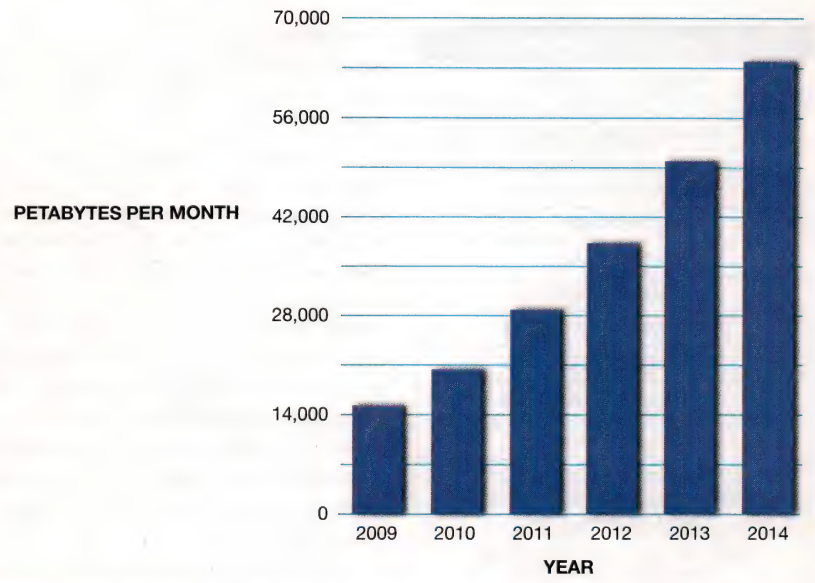

Figure 1 Video and wireless services will dominate the substantial network-data-traffic growth in coming years.

FLOW-DEFINITION FIELDS
INGRESS INTERFACE
ETHERNET SOURCE MEDIA-ACCESS-CONTROLLER ADDRESS
ETHERNET DESTINATION MEDIA-ACCESS-CONTROLLER ADDRESS
ETHERTYPE
VIRTUAL-LAN ID
SOURCE-IP ADDRESS
DESTINATION-IP ADDRESS
IP PROTOCOL
TRANSMISSION CONTROL PROTOCOL/USER DATAGRAM PROTOCOL SOURCE PORT
TRANSMISSION CONTROL PROTOCOL/USER DATAGRAM PROTOCOL DESTINATION PORT
INTERNET CONTROL MESSAGE PROTOCOL TYPE/CODE

Figure 2 A variety of packet information is available for identifying data common to a flow.

does the spatial disassociation of packets within a flow.

FLOW CONSIDERATIONS

Developers build general-purpose multicore x86 and MIPS CPUs with on-chip hierarchical caches to hide memory latencies and increase performance. For maximum cache usage and efficiency, data and instructions should reside in the on-chip cache memories for rapid access. If the cache memory does not store relevant information, the general-purpose processor instead must access external DRAM. Because external-memory bandwidths are significantly slower than those on the CPU, the processor effectively downshifts to the speed of the external memory when it is off cache.

An LRU (least recently used) algorithm typically evicts on-CPU cache memories' data. Cache-based architectures require that data flows physically reside tightly together in space, time, or both to ensure a high cache hit rate. Therefore, general-purpose multicore CPUs are ineffective at data-plane processing for networking applications because data in these types of applications is rarely spatially or temporally associative. In enterprise and carrier networks, temporal disassociation of packets is evident at all data rates, and traffic is increasingly spatially dissociative as throughputs increase. As bandwidth grows, so do the number of unique flows, making hierarchical cache memories ineffective due to low cache hit rates.

One potential approach to this architectural issue would be to simply continue increasing memory caches' sizes, but cache capacity is not keeping pace with the requirements for stateful flow processing at 10 Gbps and higher speeds. Internal processor cache memories are typically orders of magnitude smaller than external memories. For example, the Intel Xeon 5640 processor has 12 Mbytes of cache, compared with the multiple gigabytes of external memory that you can attach to these processors.

Conservatively assuming that a system requires 0.5 kbyte of memory to maintain state information for a single flow implies that, to support stateful analysis of 1 million flows, a processor's cache memory would need to be significantly larger than those that are currently available to avoid ever evicting data from the cache. Similarly, reasonably assuming a 0.5-second latency between packets in a flow on a 1-Gbps link, 500 Mbits, or almost 63 Mbytes, would have traversed the system in that short time. Assuming an average packet size of 440 bytes, more than 142,000 packets, each from different flows, would have traversed the system before the next packet within that same flow arrives. For cache-based architectures to prove effective in these demanding circumstances would require cache memories approaching 1 Gbyte in size, almost 100 times larger than those now available in multicore processors.

HITTING THE WALL

The Achilles' heel for any processing architecture is poor memory latency. When a processor's cache memories are full—that is, when they have a fully occupied memory-data structure—the CPU's cache is continuously thrashed as new packets arrive. The processor reads data once and then quickly evicts it from cache memory, causing subsequent cache misses and a low overall cache hit rate. The end result is significant memory latency when the CPU must go off-chip to get instructions and data from external memory. When this scenario occurs, the CPU stalls waiting for operations to complete, wasting as many as 200 CPU cycles per transaction for external DDR 3 SDRAM.

WHEN A PROCESSOR'S CACHE MEMORIES ARE FULL, THE CPU's CACHE IS CONTINUOUSLY THRASHED AS NEW PACKETS ARRIVE.

Purely cache-based architectures struggle to effectively handle high-packet-rate I/O traffic, security process-

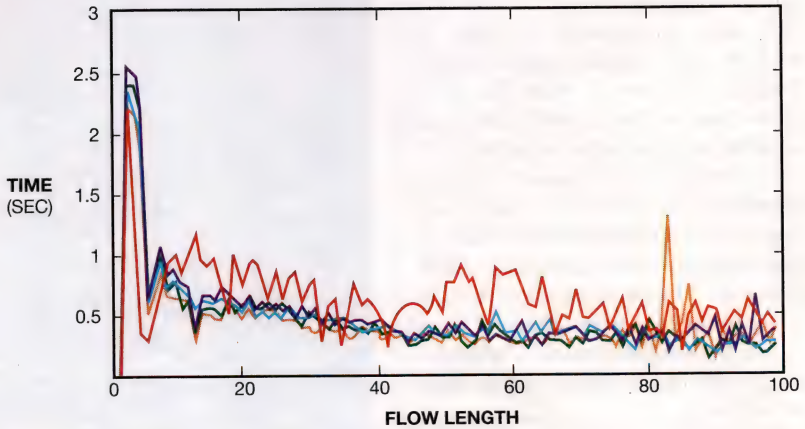

Figure 3 Packet captures of real-world network-backbone links assist in understanding the mean time between packets in a flow.

ing, and DPI at 10 Gbps and beyond. General-purpose CPUs are ideal for application and control-plane workloads, but they become a networking bottleneck in high-performance designs requiring high packet touch rates and a large number of instructions per packet over an increasing number of flows. Standard methods of hiding memory latencies, such as the use of hierarchical cache architectures, become ineffective.

Flow processors, conversely, use multiple techniques to hide memory latencies to provide more efficient memory-bandwidth usage than that of general-purpose processors. First, multiprocessing through as many as 40 independent networking- and security-optimized microengines can simultaneously process multiple independent streams of network traffic. Further, chip multithreading removes memory latency by allowing some processes to operate as other threads are waiting to complete. Using chip multithreading, memory operations are asynchronous to the processing threads. The processor structures and arranges read- and write-memory operations, effectively hiding one thread's memory references behind another thread's computations.

HETEROGENEOUS FLOWS

An effective processing architecture for intelligent network and security applications must account for cache inefficiencies that occur because of the number of instructions that the processor must apply to each packet to support stateful flow processing. Meeting these performance challenges warrants a new approach to the development of the high-performance systems that intelligent networks require. Such systems must be able to analyze traffic at all layers of the OSI (Open Systems Interconnection) model, from Layer 2, the data-link layer, to Layer 7, the application layer, and perform this intelligent processing on all traffic at sustained throughputs of 20 Gbps and higher. Achieving these goals requires specialized and varied processing elements for a specific type of workload computation.

A heterogeneous multicore architecture sets a new performance benchmark for embedded-application development through discrete processing elements for packet classification, stateful flow management, and application and control-plane processing, each with increasingly fine granularity. This architecture tightly couples network-flow-processor cores with general-purpose multicore x86 systems over a 40-Gbps virtualized PCIe (Peripheral Component Interconnect Express) datapath. This architecture can scale from low-end systems to appliances offering hundreds of gigabits per second of packet analysis, stateful flow monitoring, DPI, and application throughput, all with a common software architecture.

Accelerated designs employing this architecture can enable equipment providers to deliver high-performance, flexible systems that are more efficient than systems employing general-purpose x86 processors alone with standard NICs (network-interface cards). This architecture eliminates memory-latency problems and CPU stalls due to cache misses because the network-flow processor provides a level of preprocessing and properly structures data before transmitting it to the x86 CPU cores. Grouping traffic into flows at the network-flow-processor layer and optionally load-balancing flows across x86 cores, pinning flows to x86 destinations, or both approaches can dramatically increase the probability of cache hits because packets now arrive in a spatially and temporally associative manner. EDN

REFERENCES

1 "Cisco Visual Networking Index: Forecast and Methodology, 2009-2014," Cisco, 2010, http://bit.ly/elly1w.

2 Bradshaw, Christopher, "The Effect of Scaling Network Resources on Flows in an Internet Backbone," 2010.

AUTHOR'S BIOGRAPHY

Daniel Proch is director of product management at Netronome, where he is responsible for the company's line of network-flow-engine acceleration cards and flow-management software. He has 14 years of experience in networking and telecommunications, spanning product management, chief-technology-office positions, strategic planning, engineering, and technical support.

PICO

PROVEN Critical Component Integrity

DC-DC Converters

2V to 10,000 VDC Output

2 to 5000 VDC Outputs
Ultra Miniature
Surface Mount and Thru Hole
Single and Dual Isolated Outputs
Military Upgrades Available
AV/AV/SM/AVR Series

Military Applications
-40° and -55° to +85°C
Operating Temperatures
Wide Input Range
3.3 to 350 VDC outputs
in 1/2 Brick and Ful Brick
Military Upgrades Available
LF/LM/FM Series

Military Components
-55° to +85°C Operating Temp
Military Environmental
Screening Available
3.3 to 500 VDC Outputs
M/MV Series
Isolated - Regulated

100 to 10,000 VDC Output
Proportional Control Up to 10 Watts
VV Series

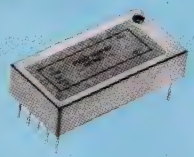

Programmable to
6000 VDC Output
HVP Series

HIGH POWER-Wide Input Range
2 to 350 VDC Output
Isolated to 300 Watts
in 1/2 and Full Brick
LP/HP/XP Series

Wide Input Range
8 to 60 VDC Input
2 to 100 VDC Output
Isolated-Regulated
OR/IR/JR/KR Series
2-20 Watts

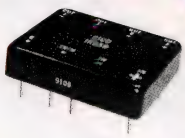

36 to 170 VDC Inputs
Terminal Strips - Thru Hole
3.3 to 48 VDC Outputs
Single and Dual Output
LV/HV Series

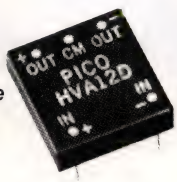

5 to 500 VDC Output
to 50 Watts
Wide Input Range
Isolated
Regulated
QP Series

Also Ac-DC single and 3 Phase.
Power Factor Corrected. 3 Watts to
2000 Watts Models.

Call Toll Free:
800-431-1064

Fax:
914-738-8225

ISO 9001:2000
ANAB ACCREDITED BSI
AS 9100 REV. B

PICO
Electronics. Inc.

143 Sparks Ave., Pelham, NY 10803

www.picoelectronics.com

Complete Listing of Entire Product Line
DC-DC Converters • AC-DC Inductors • Transformers
and Inductors. E Mail: info@picoelectronics.com


169e0
-703e7
-169e0
232e-5 453e-5 674e-5 895e-5 112e-4 134e-4 156e-4 178e-4

COMBINING FASTER SOLVERS WITH TEST BENCHES THAT LINK ANALOG TO DIGITAL TO SYSTEM-LEVEL MODELS PREVENTS DESIGN ERRORS.

COMPLETE IC SIMULATION REQUIRES A FULL TOOLBOX OF HARDWARE AND SOFTWARE

BY MIKE DEMLER • TECHNICAL EDITOR

Discovering a design error after you send your average chip to manufacturing can be costly and embarrassing, not to mention hazardous to your career. However, as Intel recently discovered, finding a design error late in a complex and widely anticipated processor SOC (system on chip) is costly. In Intel's case with its Sandy Bridge chip, the design error cost as much as $1 billion (**Reference 1**). To avoid missing design errors, you and your team must as thoroughly as possible simulate your design under the conditions that the chip will see in your customer's application. How do you simulate a complex SOC, containing perhaps 1 billion transistors, full of a combination of large digital blocks, high-performance RF and analog/mixed-signal functions, complex power management, embedded software, and gigabit I/Os? And what do you do about reliability and aging effects? Will your chip degrade over time, and how do you ensure that it will be immune to environmental problems, such as static discharge?

As the simulation challenges grow, the approaches to them also become more complex. It now takes a diverse set of simulators—for analog, digital, transistor to system level, hardware/software co-verification, chip-package interface, and more. Simulation alone is not enough, however. You also need debuggers and self-checking test benches to handle the sheer volume of data that simulation will generate. Your toolbox just got heavier.

Fortunately, although 100% assurance is usually impossible, an assortment of new tools on the market can help you gain a higher degree of confidence. These tools run the gamut from fast circuit simulation to real-world emulation of system applications with embedded software and data I/O. A one-size-fits-all approach is not available, but there is a lot to choose from.

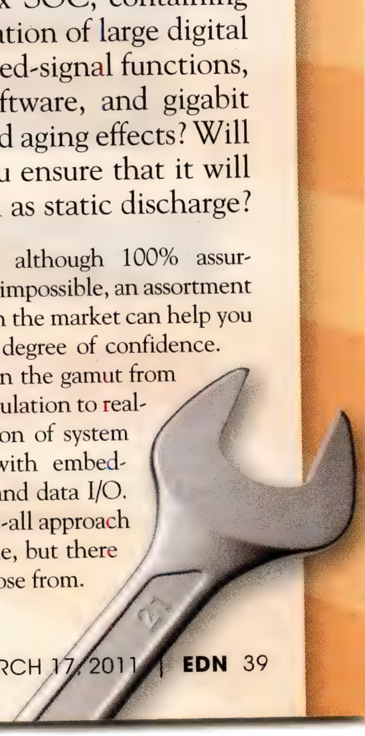

PHOTO-ILLUSTRATION BY TIM BURNS. TOOLBOX: SERGEY KULIKOV/ISTOCKPHOTO.COM; CHIPS: CRAFTVISION/ISTOCKPHOTO.COM; PROBES: IGOR SMICHKOV/ISTOCKPHOTO.COM; WRENCH: MAI FRAPASO/ISTOCKPHOTO.COM

SPEEDING SIMULATION

There is not much in the way of new analog simulators, according to KT Moore, vice president of simulation technology for Magma Design Automation's custom-design business unit. Gary Smith, the founder and chief analyst at Gary Smith EDA, agrees. "There was some hope that parallel computing would be a big thing," he says, "but it hits the Amdahl's Law barrier of about four processors," referring to the idea that Gene Amdahl stated in 1967 that, even when the fraction of serial work in a given problem is small—say, s—the maximum speedup you can obtain from even an infinite number of parallel processors is only 1/s.

Moore also points to parallel processing, saying that Magma has overcome the single-matrix limitation of Spice, which peaks at about 25 thousand devices, by developing methods to scale the problem across a network with the company's FineSim simulator. This approach differs from that of FastSpice simulators, which typically break the matrix of large circuits into smaller partitions with simplified device models, potentially introducing errors and loss of accuracy. Magma's approach is different, he explains, because the company has figured out how to maintain full BSIM (Berkeley Short Channel IGFET Model) accuracy—equivalent to conventional Spice simulators—with circuits having as many as 1 million components. FineSim can handle fully extracted postlayout simulations by distributing analysis of a single matrix across eight to 16 CPUs in two to four machines, according to Magma. This approach has changed the game, Moore adds.

Magma's customers are looking for a more digital-centric verification method, says Moore. The ability to simulate a large design is not enough on its own because users can't afford to look at all the waveforms. "You want to look at signals only when there is a problem and just have a pass/fail indicator otherwise." It is possible to have conditional checking, but debugging must be more stringent. Otherwise, engineers will be unsure about whether they have missed something. "We need to think about what problem customers are trying to solve, as well as performance and accuracy," he explains.

Simon Young, product-marketing manager at Berkeley Design Automation, also finds that capacity limitations have forced customers to use FastSpice but offers a different opinion on what is new in analog simulation. You must account for new device behaviors, such as noise, in 28-nm processes, he says. Berkeley Design has addressed this problem with the transient-noise-analysis feature in its AFS (Analog FastSpice) platform. The company claims that AFS, which has capacities as high as 10 million elements, is 10 times faster on a single-core CPU than competing versions of Spice. Semiconductor foundry TSMC (Taiwan Semiconductor Manufacturing Co) has certified AFS for its 28-nm, low-power-process design flow, and Berkeley Design is also working with GlobalFoundries, according to Young.

Engineers have been trying for many years to solve the problem of how to accurately simulate ESD (electrostatic-discharge)-protection devices. The PathFinder tool from Apache Design Solutions addresses this challenge. According to Andrew Yang, PhD, chief executive officer at the company, traditional transient solvers cannot handle the "snap-back effect" in ESD structures. In snap-back, when the voltage exceeds the trigger voltage, the IV (current-to-voltage) characteristic of the device snaps back, allowing the same or a higher amount of current to flow but at a significantly lower voltage. Negative resistance, Yang explains, causes convergence problems for traditional Spice simulators. PathFinder relies on the proprietary eSim nonlinear transient simulator, which works with standard Spice-device models. PathFinder has built-in extraction for power- and

AT A GLANCE

- Parallel processing and specialized solvers are speeding analog simulation.
- New industry standards for testbench creation mean more choices in digital simulators.
- Synthesis of test vectors improves productivity and eliminates redundant simulations.
- A faster simulator is not enough; debugging and linking abstraction layers are also necessary.
- Hardware acceleration is staging a comeback as software struggles with design complexity.
- Standards are in development to lower the barrier between analog and digital simulation.

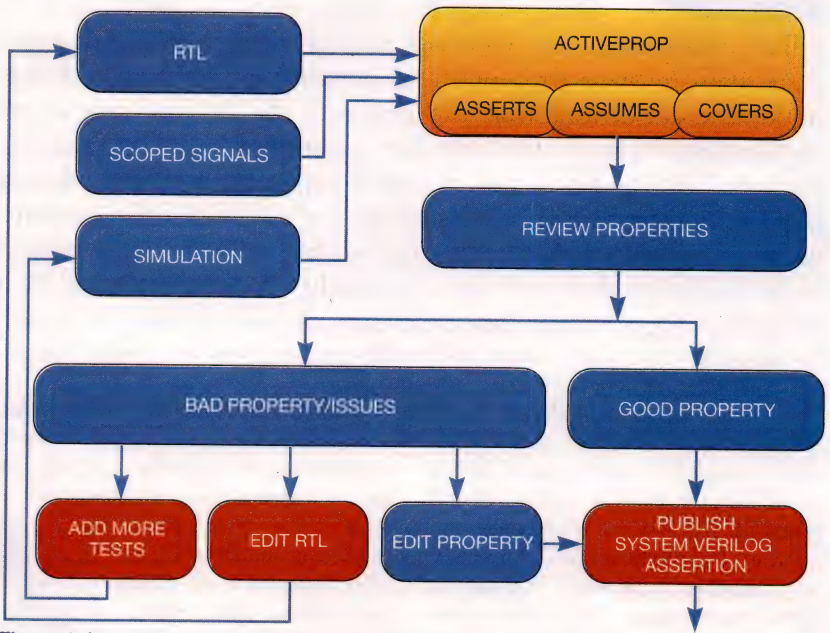

Figure 1 Jasper's ActiveProp analyzes and extracts critical expressions from your RTL in combination with simulation data to generate SV constraints, assertions, and coverage properties that it ranks and merges to reduce the number of candidates for review.

"I need the fastest DMM they make."

They don't make them any faster.

Nobody but Agilent makes a range of DMMs this fast, this accurate, or this reliable. Up to 1000 times more readings per second than the nearest competitor's, and far easier to use, you'll rip through tests in a fraction of the time. It's what you'd expect from the leader in DMM technology.

DMM	Digits	DC Accuracy	Max Readings	Function/Range Changes	IO
34405A	5 ½	0.0250%	19 / sec	0.2 sec	USB
34401A	6 ½	0.0035%	1,000 / sec	.02 sec	GPIB, RS-232
34410A	6 ½	0.0030%	10,000 / sec	2.6 ms	GPIB, USB, LAN (LXI)
34411A/ L4411A	6 ½	0.0030%	50,000 / sec	2.6 ms	GPIB, USB, LAN (LXI)
34420A	7 ½	0.0030%	250 / sec	.02 sec	GPIB, RS-232
3458A	8 ½	0.0008%	100,000 / sec	3.0 ms	GPIB

Agilent and our Distributor Network

Right Instrument.
Right Expertise.
Delivered Right Now.

METRICTEST™
6,000 instruments. One source.

866-436-0887
www.metrictest.com/agilent

Download the latest measurement brief and tips
www.metrictest.com/agilent/dmm.jsp

Agilent Technologies

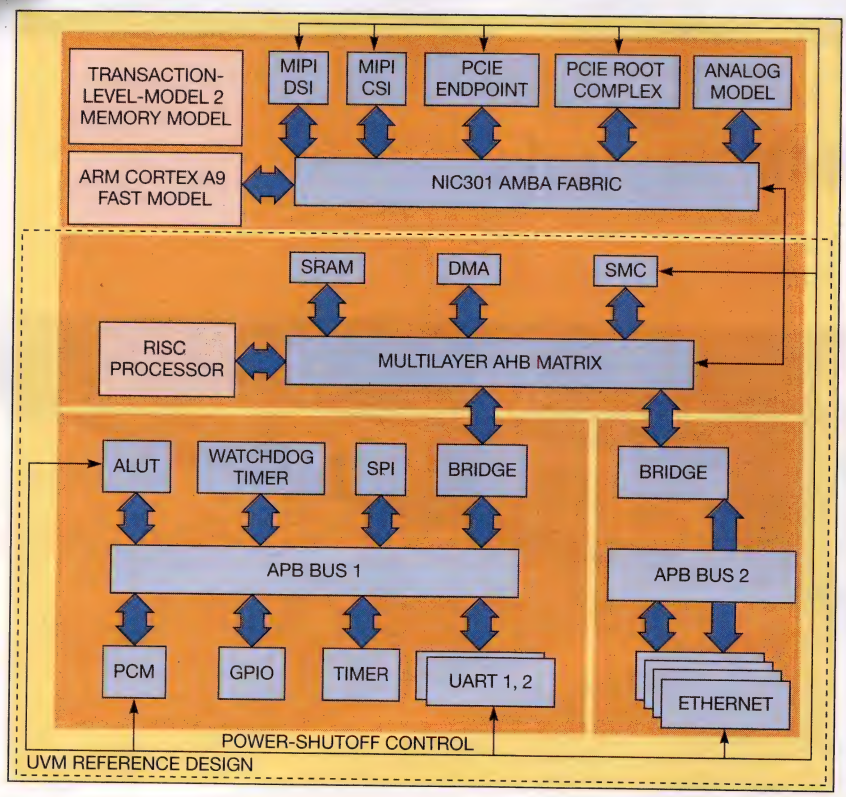

AHB: ADVANCED HIGH-PERFORMANCE BUS
ALUT: ADDRESS-LOOK-UP TABLE
AMBA: ADVANCED MICROCONTROLLER-BUS ARCHITECTURE
APB: ADVANCED PERIPHERAL BUS
CSI: CAMERA SERIAL INTERFACE
DMA: DIRECT-MEMORY ACCESS
DSI: DIGITAL SERIAL INTERFACE
GPIO: GENERAL-PURPOSE INPUT/OUTPUT
MIPI: MOBILE-INDUSTRY PROCESSOR INTERFACE
PCM: PULSE-CODE MODULATION
SMC: SMART MEMORY CONTROLLER
SPI: SERIAL-PERIPHERAL INTERFACE
SRAM: STATIC RANDOM-ACCESS MEMORY
UART: UNIVERSAL ASYNCHRONOUS RECEIVER/TRANSMITTER

Figure 2 Cadence's Incisive verification kit includes a prototype 1.5 million-gate design for an Ethernet-controller SOC as an aid to learn functional-verification concepts, along with a set of self-paced workshops.

ground-bus-RLC (resistor/inductor/capacitor) elements and substrate-package parasitics. According to Yang, eight large semiconductor companies have adopted PathFinder, which can perform layout-based analysis of circuits with more than 1 million elements for ESD events, such as human-body model, machine model, and charged-device model. TSMC's Reference Flow 11.0 also includes PathFinder.

TEST-BENCH STANDARDS

According to Steve Bailey, director of marketing for verification and test at Mentor Graphics, as the market for simulation has matured, focus has shifted to simulation tools that can help engineers to more quickly complete verification. Although techniques such as parallel processing, multicore solvers, and multithreading improve simulator performance, they aren't silver bullets because they don't ease the task of finding an optimal match of solvers with partitions of a design.

In a recent survey of users, Mentor Graphics found a higher-than-expected adoption of the SV (System Verilog) HVL (hardware-verification language), with three-fourths of survey participants saying that they either were using or planned to use SV for creating test benches. One of the issues limiting SV's adoption had been a lack of compatibility among EDA vendors' tools. In choosing from a variety of available SV-based verification methods, you must ensure that your simulator is compatible with your test-bench libraries.

Current industry efforts to standardize a base library of functions for digital verification are addressing the portability issue. According to Bailey, the development of the new Accellera UVM (Universal Verification Methodology) Version 1.0 EA (early-adopter) standard represents a maturation of SV techniques that companies first introduced in predecessor products, starting with Mentor's AVM (advanced verification methodology), which OVM (open verification methodology), a collaboration between Mentor and Cadence, followed. Elements of the Synopsys-developed VMM (*Verification Methodology Manual*) are also part of the new standard.

Michael Sanie, director of verification marketing at Synopsys, points out that Synopsys has participated in Accellera's efforts for development of UVM, with representation by VMM expert and Synopsys Fellow Janick Bergeron. UVM incorporates VMM features, including the RAL (register-abstraction-layer) application package for memories and memory-mapped

Simple, High-Current LDOs for Powering Next Generation FPGAs, DSPs and ASICs

Single Input Supply, Ultra-Low V_{IN}/V_{OUT} LDOs

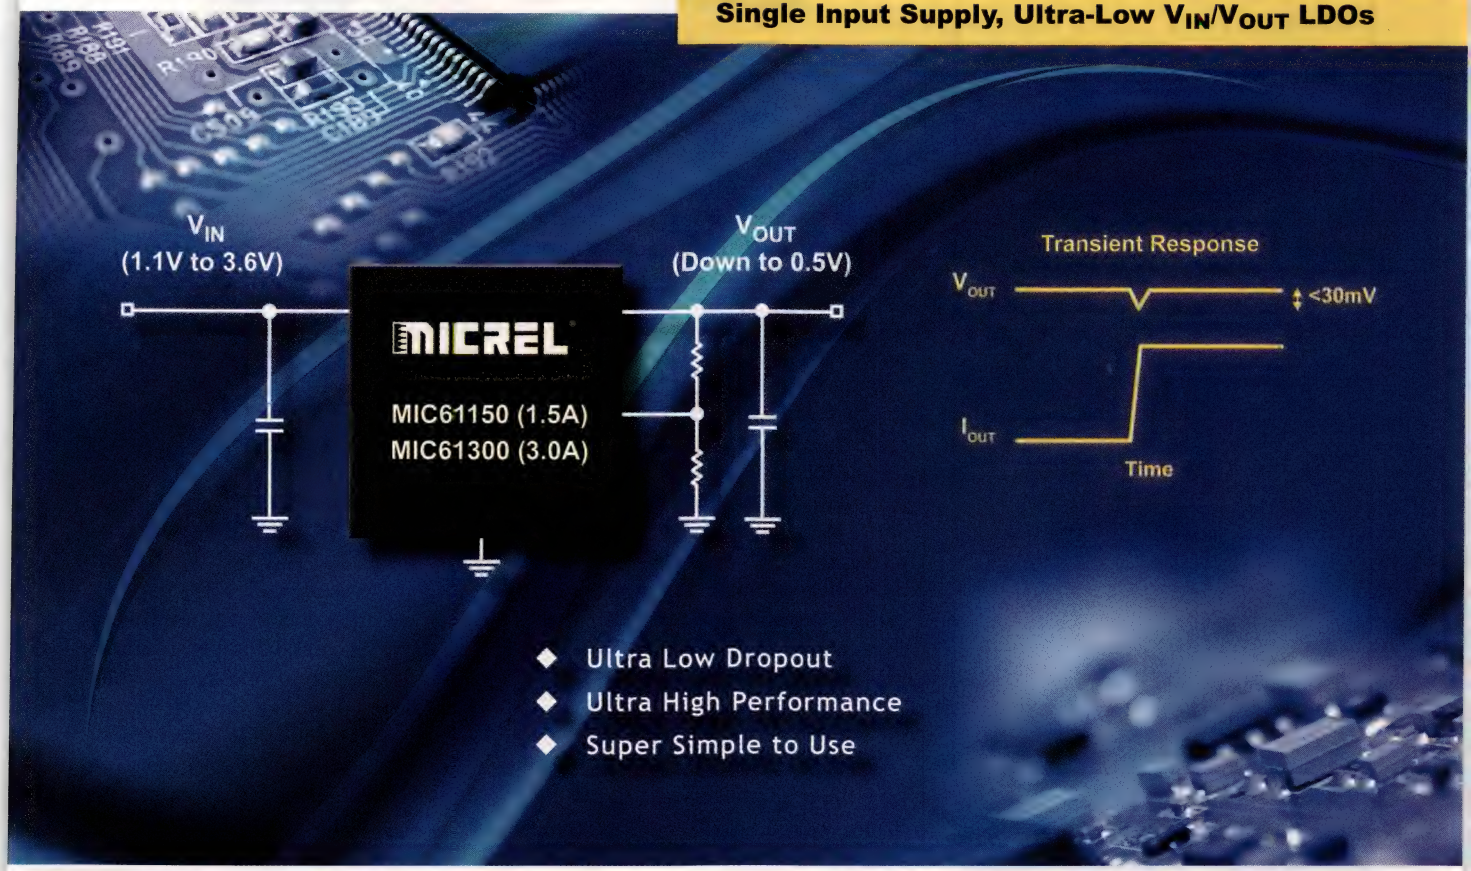

Powering today's FPGAs, ASICs and DSPs involves the challenge of supplying high currents in very low input-output voltage conditions. This is where Micrel's new family of ultra-low V_{IN}/V_{OUT} LDOs excels. These new Micrel LDOs offer the convenience of adjustable output voltages as low as 0.5V with input voltages as low as 1.1V, while supplying up to 3A of output current.

To further enhance flexibility and reduce overall cost for the power supply designer, Micrel has eliminated the additional bias supply! Unlike other solutions, the MIC61150 and MIC61300 do not require an additional bias supply and hence offer clean and simple single input supply operation.

For more information, contact your local Micrel sales representative or visit Micrel at: www.micrel.com/ad/mic61xxx.

Key Applications

- FPGA/DSP Power Supplies
- ASIC/Microprocessor Power Supplies
- Telecom/Networking Cards
- Point-of-load Applications

MICREL®

Innovation Through Technology®

www.micrel.com

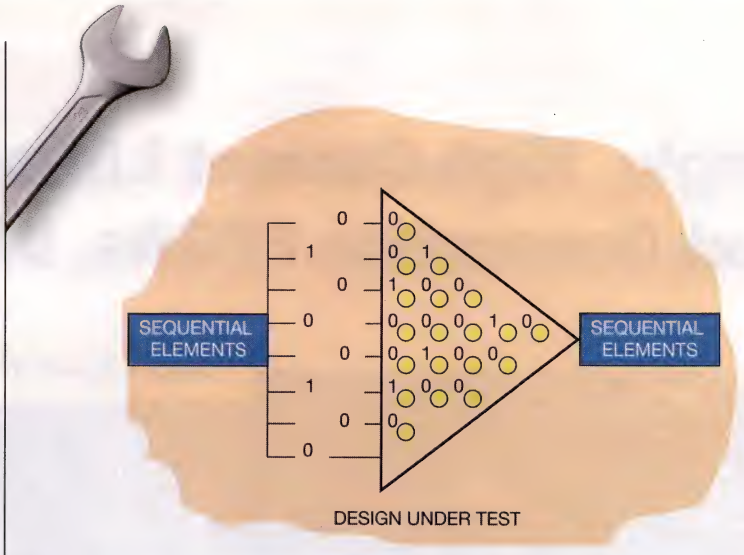

Figure 3 Combinational signal access in Eve's ZeBu platform allows you to probe the values of any RTL net in a cone of logic between registers.

registers. Synopsys will provide native support for UVM 1.0 EA in May.

Sanie also notes that users of Synopsys' VCS HDL simulator can use OVM, so they need not wait for UVM to gain test-bench portability. He believes that the method of writing base classes for verification libraries is not the key differentiator among vendors' tools. Vendors will continue to compete to establish an advantage in performance of the simulation engines and in debugging tools.

AUTOMATE TEST GENERATION

Verification engineers can specify how the properties of a design will undergo testing with SVA (System Verilog-assertion) statements, define constraints for random stimuli, and determine how to check coverage. The ActiveProp tool from Jasper Design Automation automatically synthesizes these properties from RTL (register-transfer-level) representations of your design, or you can reuse simulation and test-bench waveforms in VCD (value-change-dump) or FSDB (fast-signal-database) formats.

According to Holly Stump, vice president of marketing at Jasper, ActiveProp can jump-start your ASV (assertion-based verification) by automatically generating assertions for formal analysis, simulation, or emulation. ActiveProp can work as a stand-alone tool (**Figure 1**), or you can link it at runtime to simulators through an API (application-programming interface). By performing multiple simulation runs, you can use the results to further refine the intelligence of generated properties. The first release of ActiveProp supports Synopsys simulators.

When you submit your design to ActiveProp, it analyzes and extracts critical expressions from your RTL to guide, simplify, and merge the generated properties. To initiate property synthesis from RTL, you scope your signals by identifying RTL-block signals and points of interest. You can use this information in combination with simulation data as input to ActiveProp, which then generates ranked and merged properties to reduce the number of candidates for review. Output is in the form of SV constraints, assertions, and coverages, along with reports in human-readable format.

Adoption of techniques such as constrained random stimuli in the UVM standard can help to improve verification coverage over manual, or directed, test-pattern generation, but it is also critical to eliminate redundancy in the generated test patterns. To address this problem, Mentor offers the inFact intelligent test-bench-automation tool that you can use to build test-bench components. So-called rule graphs control the activity of these components. With inFact, you can define rules from the specifications of the device you are verifying, such as for an industry-standard bus interface, or by applying constraints from your verification plan. When you run the simulation, the rule graphs interact with the tools and apply intelligent algorithms to more efficiently achieve the required functional coverage. The inFact tool can distribute simulation to multi-

Introducing ZYNQ, the new element in processing. Finally, the processor comes together with the FPGA in a fully extensible processing platform called Zynq.™ More intuitive to program in the way you already know. Fully customizable to your requirements. Faster to implement and get to market. As a software engineer, if you know ARM® Cortex,™ you already know Zynq. And if you know Xilinx, you already know this is innovation you can count on. Visit us at www.xilinx.com

XILINX
ZYNQ™

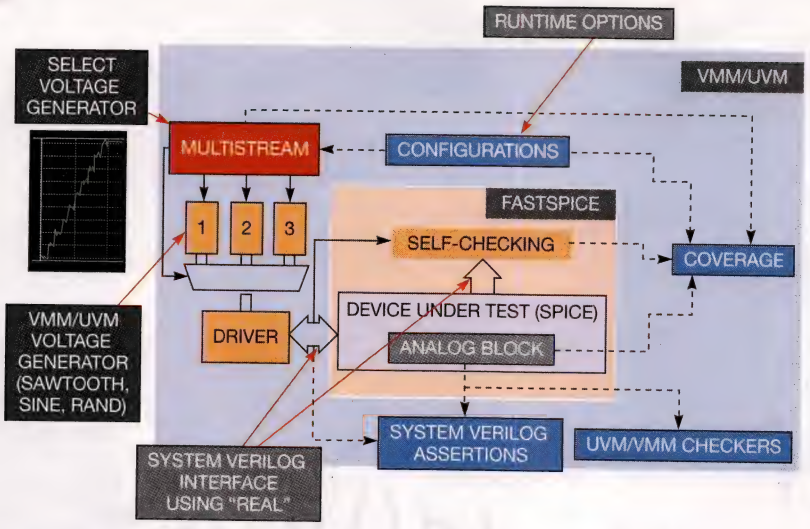

Figure 4 Synopsys' AMS test bench extends VMM and UVM test benches to enable the VCS HDL simulator to read or write voltages or currents to or from the XA and NanoSim FastSpice engines in CustomSim.

ple CPUs and performs load balancing. According to Mentor's Bailey, inFact's graph-based-testing customers have experienced as much as a 100-times speedup in the time it takes to achieve coverage, and the coverage percentage has also increased.

INTEGRATE TEST BENCHES

According to Tom Anderson, marketing manager for verification products at Cadence Design Systems, engine-level performance alone is simply not enough to solve the verification problem. The complexity of design has caused fragmentation of the verification process into niches, such as low power, RTL simulation, mixed signal, and formal analysis. Adam Sherer, Cadence's verification-product-management director, echoes that opinion, saying that, although simulation performance is key, it isn't enough. "As elaboration times grow for large SOCs, both faster elaboration and incremental elaboration are also necessary," he says. "Customers need faster waveform and coverage database access and smaller database size."

With Cadence's latest 10.2 release of the Incisive verification tools, you can link design and power intent into one verification plan and combine formal models with digital- and mixed-signal abstractions in the MDV (metric-driven-verification) cockpit. With MDV, verification engineers can merge coverage data from formal analysis with

results from simulation engines into a unified SOC-level verification plan. By linking formal and dynamic analysis in MDV, you can use the formal tools to flag points in your design that are unreachable, indicating that you shouldn't try to simulate them. Incisive can also automatically generate assertions and coverage points for simulation or formal analysis so that it is easier to define which simulators to use to execute each part of the verification plan.

Real-number models support the interface of digital- to mixed-signal behavior in Incisive. In Incisive 10.2, Cadence has optimized simulator performance for Accellera's new UVM standard for test-bench development. Incisive integrates UVM components as native objects in simulation, with the ability to monitor transactions in the waveform viewer. Cadence has extended OVM and UVM for the 'e' verification language and has facilitated the use of multiple languages in test benches. It supports all IEEE-standard languages, including Verilog, VHDL (Very High Speed Integrated Circuit Hardware Description Language), PSL (property-specification language), 'e,' System-C, and System Verilog. To improve performance for large numbers of regression tests, Incisive's advanced Specman option provides reseeding and dynamic loading of 'e'-based tests, multicore 'e'-code compilation, and the ability to shorten debugging time by mixing interpreted and compiled code.

Users of the Incisive Enterprise Simulator also receive the Incisive verification kit, a 1.5 million-gate design for a prototype Ethernet-controller SOC that acts as a demonstration vehicle for the functional-verification concepts (**Figure 2**). With the verification kit, you can learn the new technologies and methods through a set of self-paced workshops. The UVM reference flow is available as open source in the contribution area on the UVM World Web site.

SPEEDING SIMULATION

Lauro Rizzatti, general manager of Eve USA, points to EDAC (Electronic Design Automation Consortium) data that shows that emulation and HDL (hardware-description-language) accel-

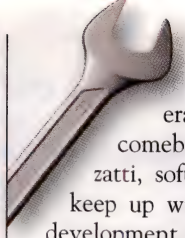

erators are making a comeback. According to Rizzatti, software simulation can't keep up with embedded-software development and whole-SOC verification. Although revenue in the EDA-simulation segment peaked at $532 million in 2007 and declined on an annual basis to $323 million last year, emulation has been relatively steady in the $140 million to $160 million range, accounting for a larger percentage of the overall verification market. Hardware acceleration was big 10 years ago with CPU or GPU (graphics-processing-unit) companies, says Rizzatti, but usage is now growing in the wireless, multimedia, networking, and image-processing markets. Hardware accelerators work with test benches for industry-standard Verilog-, SV-, and VHDL-based software simulators, letting you run simulation at millions of cycles per second.

The cost of acceleration hardware has historically limited adoption, an issue that Eve has addressed by building its systems with Xilinx Virtex-8 FPGAs versus ASICs. The trade-off is that the FPGA approach has more limited visibility of results, but Eve has now extended debugging capability by adding CSA (combinational signal access) to its ZeBu (zero-bug) platform. CSA allows you to use dynamic probing to see the values of any RTL net in a cone of logic between registers or memory in your design (**Figure 3**). You can view results in standard VCD or FSDB formats. You can also do offline analysis with CSA by generating waveforms after your emulation run on a PC or Linux workstation.

Mentor Graphics' Bailey has also seen an increase in the emulation- and simulation-acceleration market. Hardware acceleration is necessary for designs such as multimedia processors that require test sequences that are too long to run using software simulation alone. Mentor's Veloce system can handle large SOC-simulation-acceleration tasks and enables software/hardware co-verification. You can use Veloce as an in-circuit emulator, connecting the model of your SOC to a physical-target design-verification environment to run embedded software and use real-world data to exercise the DUT (device under test).

REMOVING THE BARRIERS

The boolean nature of digital circuitry lends itself to a level of automation and abstraction that does not exist in the analog world. AMS (analog/mixed-signal) modeling languages, such as Verilog-AMS, are now more than 10 years old, but Spice and FastSpice remain the workhorses for simulation and analysis of analog behavior. Cosimulation, linking Spice or FastSpice tools with an HDL simulator, can be an effective way to bridge the gap when you need to verify mixed-signal behavior that involves analysis of analog-to-digital interactions.

All SOCs incorporate some form of analog behavior that you must simulate and verify, whether in power-management and clocking circuits or in more advanced signal-processing blocks and ADCs. With SV gaining popularity for SOC test benches, the need to cover both analog and digital behavior has led to a number of efforts to extend the SV language for AMS verification.

Synopsys' AMS test-bench technology extends VMM and UVM with new constructs for analog blocks (**Figure 4**). The company added proprietary SV functions to its VCS HDL simulator to allow SV test benches to check designs that you cosimulate with Synopsys CustomSim. According to Bradley Geden, CustomSim product-marketing manager at Synopsys, the AMS test-bench integration supports cosimulation of the VCS HDL simulator with the XA and NanoSim FastSpice engines within CustomSim.

FOR MORE INFORMATION

Accellera
www.accellera.org/home

Apache Design Solutions
www.apache-da.com

Berkeley Design Automation
www.berkeley-da.com

Cadence Design Systems
www.cadence.com/us/pages/default.aspx

Eve
www.eve-team.com

Freescale Semiconductor
www.freescale.com

Gary Smith EDA
www.garysmitheda.com

Intel
www.intel.com

Jasper Design Automation
www.jasper-da.com

Magma Design Automation
www.magma-da.com

Mentor Graphics
www.mentor.com

Synopsys
www.synopsys.com/home.aspx

Taiwan Semiconductor Manufacturing Co
www.tsmc.com

UVM World
www.uvmworld.org

Xilinx
www.xilinx.com

With the AMS test-bench technology, you can use VCS to read or write voltages or current to or from the analog simulator and convert the values to real numbers. This approach trades off performance versus accuracy compared with a discrete, piecewise-constant real-number model because cosimulation of the actual circuit derives the analog values in the AMS test bench at runtime. With the AMS test bench, you can monitor asynchronous interactions that better adapt to the unclocked nature of most analog behavior.

Components of the VCS-AMS test bench include a base class of waveform generators and checkers. Stimuli are fully customizable, and the tool provides sine-, sawtooth-, and square-wave generators. Prebuilt checkers measure common analog behaviors, such as thresholds and stability, frequency, and slew-rate measurements. With the VCS-AMS test bench, you can apply SV-constrained random-stimulus techniques to analog signals, and you can model mixed-signal transactions as an SV class. You can also apply SV monitors to analog signals to check with SVAs. By using VMM as the top-level test bench, your verification can include more complex test sequences that you can exercise without separate simulations.

To perform asynchronous sampling of analog signals, you can apply Synopsys' adaptation of the Verilog-AMS cross function: always @(snps_cross()). As with Verilog-AMS, you must minimize the sampling rate of signals to avoid too severely affecting simulation performance. You can group signals together in SV-interface functions. By instantiating your DUT in an AMS test bench, VCS will automatically perform e2r (electrical-to-real) and r2e (real-to-electrical) transformations. You can access nodes anywhere in the hierarchy of your Spice or Verilog-AMS circuit, to get or force voltages or port currents.

To apply the SVA feature on analog signals, you can create both synchronous and asynchronous monitors. To create a synchronous monitor, you use a clock from the digital simulator to initiate a test. For an asynchronous monitor, Synopsys has added a system function that allows the analog simulator to initiate the test.

AMS VERIFICATION

According to Scott Little, AMS verification engineer at Freescale Semiconductor and chairman of Accellera's Verilog-AMS Assertions Subcommittee, a need exists for an AMS equivalent to UVM to bridge the gap between analog and digital verification. The subcommittee is attempting to define extensions to SV that will allow for the intermingling of continuous, real-time behaviors with the discrete-time nature of digital standards, such as UVM. "We are trying to add more rigor to the verification engineer's toolbox," he says. "When a digital-verification engineer bumps up against the analog boundary, there are not a lot of tools to address the problem. Standards are not so strong in the analog world; they are much better in digital."

Designers must increasingly rely on digital control of analog circuits to meet performance requirements in SOCs. Engineers must also combine analog and digital functions to manage the complex low-power modes that enable packing more functions into a die without exceeding package and supply limitations. These two factors are driving the Accellera subcommittee's efforts to define a language for AMS assertions. Researchers must address some theoretical roadblocks that remain for extending a number of SVA concepts to analog, but the committee will focus on a specified, certifiable subset, says Little. The group's direction is to add SVA to Verilog-AMS so that engineers can write mixed-signal assertions in a digital block. The subcommittee hopes to incorporate these new functions on the next revision to the Verilog-AMS specification, which should emerge this year. EDN

REFERENCE

1 Demler, Mike, "Price for a new SATA I/O: $700 million. A complete AMS verification? Priceless!" *EDN*, Feb 1, 2011, http://bit.ly/ejQaPr.

BY MARTIN NUSS, PhD • VITESSE

Hardware choices for MPLS and backbone bridging

BOTH GROWING SILICON CAPABILITIES AND THE NEEDS OF SERVICE PROVIDERS WILL INFLUENCE THE CONTEST BETWEEN MPLS AND PBB-TE AT THE NETWORK EDGE.

The traditional model of evolving hardware support for more complex routing in carrier-based network equipment assumed that smaller silicon features, more optimized design techniques, and stabilized high-layer protocols would converge to allow single-chip switch devices to climb up the protocol stack. Switches that once supported only Layer 1 PHY (physical) and Layer 2 data-link functions would grow over time to include Layer 3 network functions, Layer 4 transport duties, and a mix of higher-layer presentation and application functions to support duties such as deep-packet inspection.

In a broad-brush approach, this vision still holds. The arrival of MPLS (multiprotocol label switching) to provide IP (Internet Protocol)-switching capabilities in scaled hierarchies of routers played an important role in simplifying routing for silicon. In the middle of the last decade, however, some equipment providers and carriers expressed interest in PBB (provider backbone bridging), a simple Layer 2 protocol, throwing a wrench into the just-simplified works. In some sense, this move was a replication in the public network of the bifurcation in enterprise networks between routing and bridging. Consequently, you could look for lessons from the evolution of bridging that took place in the early 1990s. As you will see, however, you cannot assess the current "debate" between PBB and MPLS as a simple replay of the LAN (local-area-network) bridging-routing dichotomy, taken to the WAN (wide-area-network) domain.

In LAN-based bridging, packets using similar protocols transfer from one subnet to another. Ethernet's shift from shared hub and bus topologies to switched hub-and-spoke topologies in the early 1990s streamlined and simplified bridging. As Layer 2 switching displaced all shared-Ethernet LANs by the end of the century, silicon providers were able to standardize switch designs and integrate multiple ports on one switch device. At first, vendors offered quad and octal switches without PHY or full MAC (medium-access-control) functions on the switch chip. As semiconductor providers turned to deep-submicron processes, however, switch chips scaled to 24 ports and more, even as they integrated MAC and PHY functions in the central switch chip.

Switch chips for the LAN could improve along another dimension, as well. Centralized routers and edge-based access devices, which had to interface with protocols other than Ethernet, added management intelligence and MPLS support to their core switch chips, thus making these devices Layer 3 switches. This process began with ASICs and FPGAs because many early routers relied on proprietary interfaces to WANs or dissimilar LANs. The victory of Ethernet and TCP (Transmission Control Protocol)/IP over ATM (asynchronous transfer mode), however, combined with the standardization of edge-routing functions in the enterprise and allowed the Layer 3 switch to become an ASSP (application-specific standard product). Easier integration of glue logic allowed vendors to integrate the Layer 3 functions into reasonably priced switches, but a Layer 3 switch always carried a price premium over its Layer 2 equivalent. The semiconductor-developer world adopted a slogan that originated in the system-level-network-equipment world: "Bridge—or switch at Layer 2—when you can and route when you must." To a certain extent, that same philosophy applies in today's service-provider WAN, although the debate between PBB and MPLS is not a simple replication of the bridging debate of the past.

The newest version of PBB, PBB-TE (PBB transport engineering), suffers from the fact that the word "bridging" remains in the technology description. The type of backbone bridging that 21st-century multigigabit backbones use is not your father's bridging, and its complexity stems from more than just the order-of-magnitude improvements in speed.

MPLS AND PBB-TE: WHAT'S THE POINT?

A natural tension exists between a connection-oriented deterministic world, in which service guarantees are a natural response to the problem of deterministic delays, and the connectionless packet world of the Internet. This tension forms the basis of the transport technologies that MPLS and PBB-TE defined. Although the early TCP/IP networks had no true nature-of-service guarantees, in practice carriers did offer such assurances in a de facto manner at higher network layers, to give a connectionless, packet-switched topology the same robustness as SONET (synchronous optical network)/SDH (synchronous digital hierarchy).

Before considering the relative merits of either PBB-TE or MPLS, it is fair to ask whether users still need any form of connection-oriented emulation, now that advances in deep-submicron silicon allow significant improvements in speed of IP-header look-ups. As TCP/IP began overtaking ATM

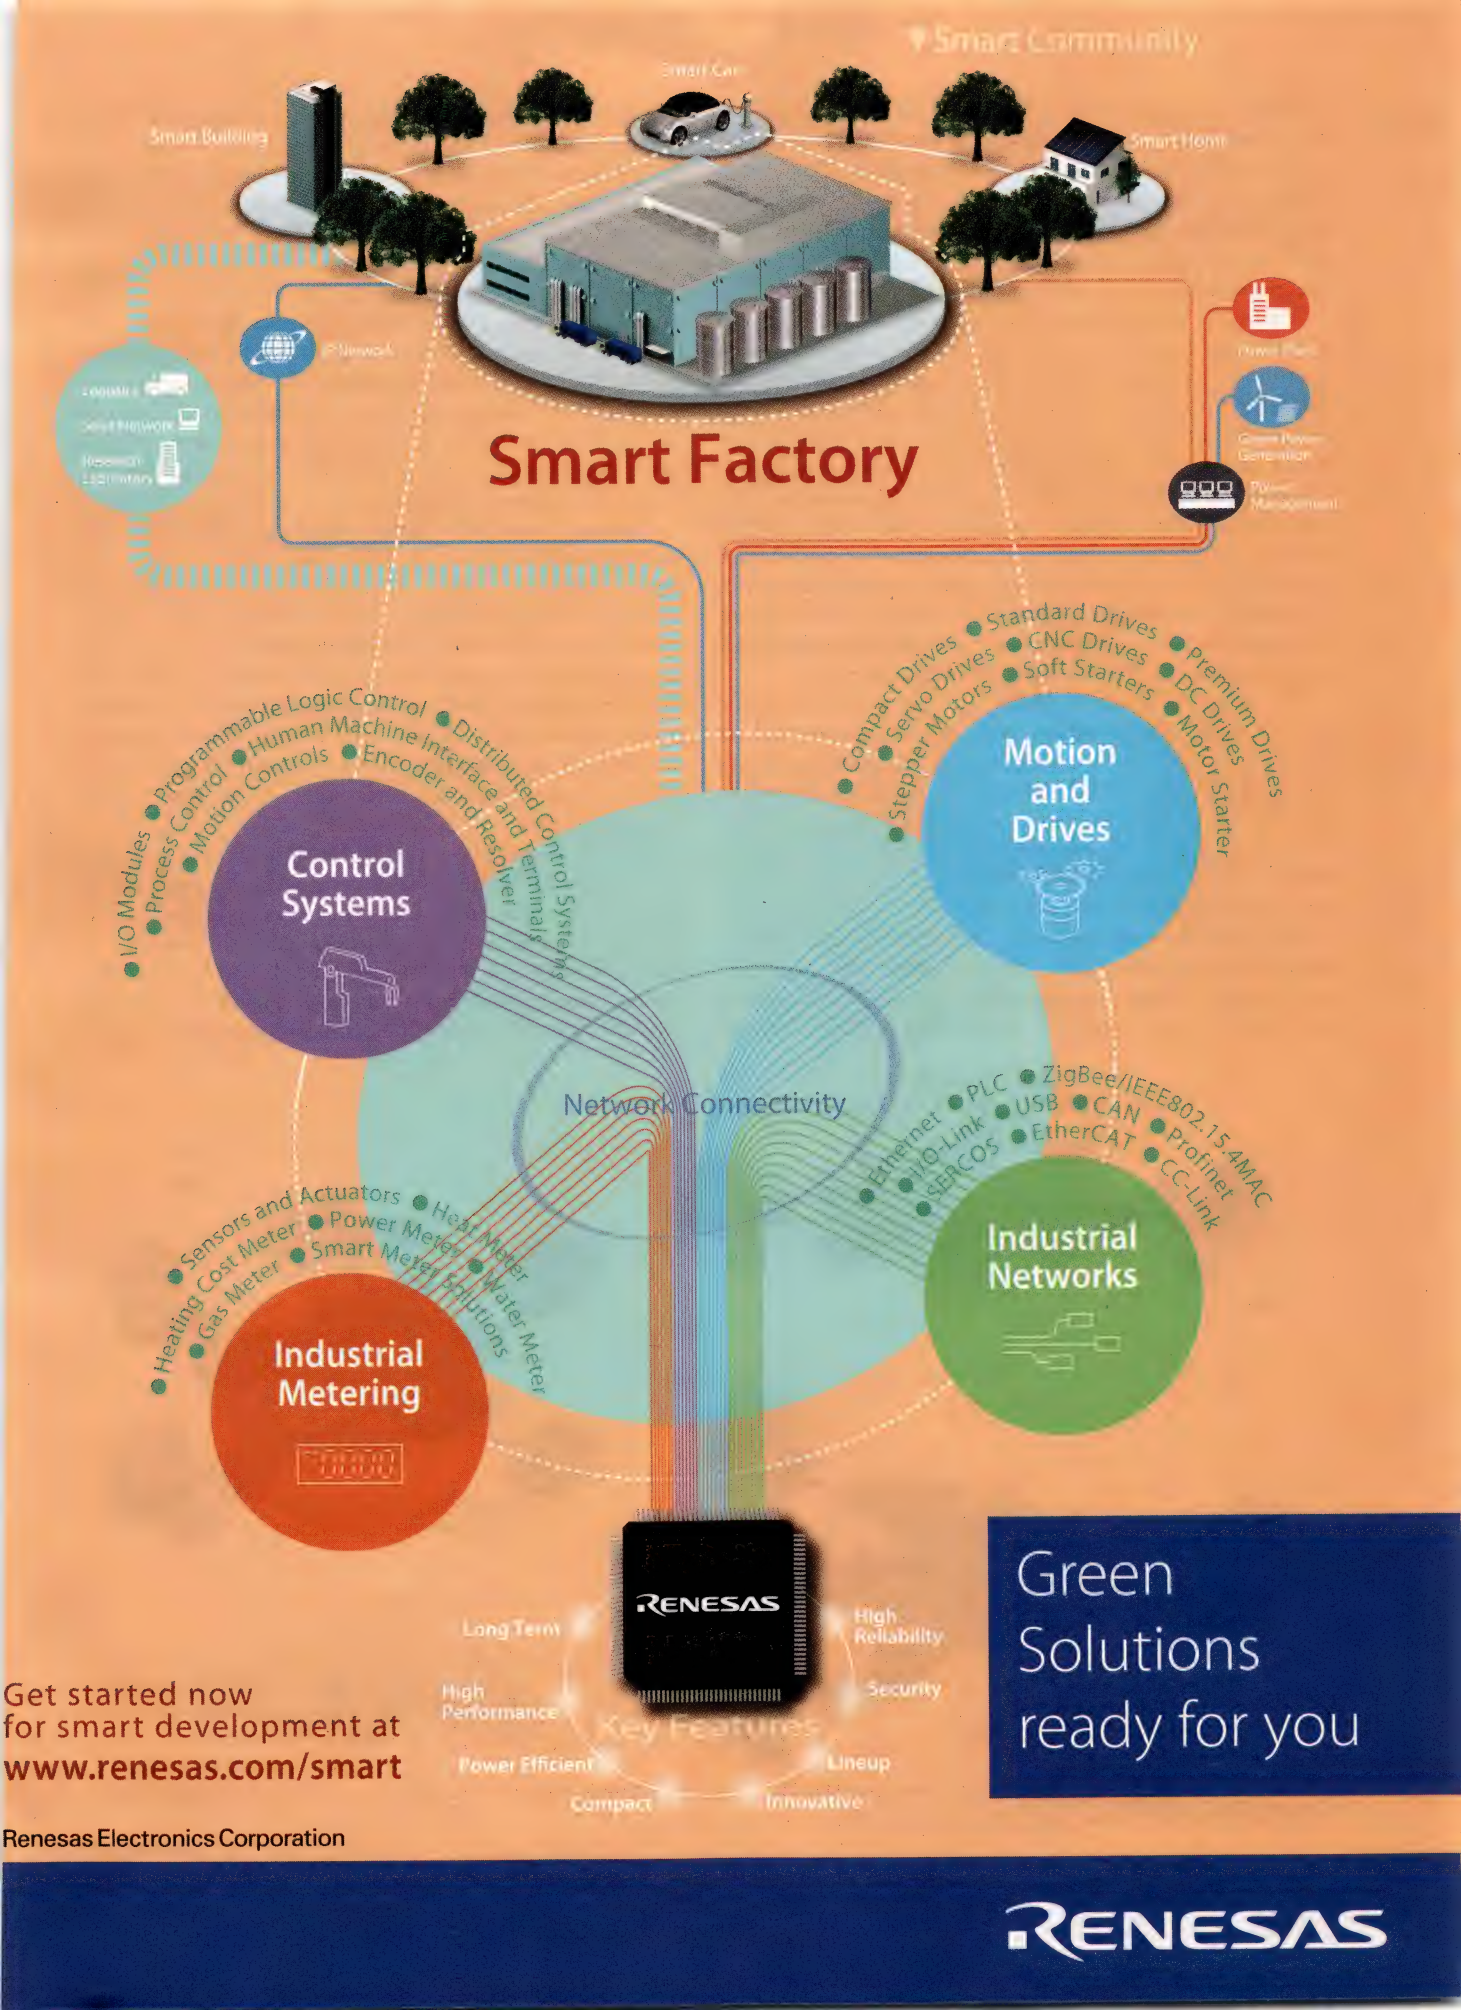
Smart Community
Smart Car
Smart Building
Smart Home
IP Network
Smart Factory
Power Plant
Green Power Generation
Power Management
Programmable Logic Control
Human Machine Interface and Terminals
Distributed Control Systems
Encoder and Resolver
Motion Controls
Process Control
I/O Modules
Control Systems
Compact Drives
Standard Drives
CNC Drives
Premium Drives
Servo Drives
Soft Starters
DC Drives
Stepper Motors
Motor Starter
Motion and Drives
Network Connectivity
Ethernet
PLC
ZigBee/IEEE802.15.4MAC
I/O-Link
USB
CAN
Profinet
SERCOS
EtherCAT
CC-Link
Industrial Networks
Sensors and Actuators
Heat Meter
Power Meter
Heating Cost Meter
Smart Meter Solutions
Water Meter
Gas Meter
Industrial Metering
RENESAS
Long Term
High Reliability
High Performance
Security
Key Features
Power Efficient
Lineup
Compact
Innovative
Green Solutions ready for you
Get started now for smart development at
www.renesas.com/smart
Renesas Electronics Corporation
RENESAS

switching in the late 1990s, developers looked for some way to provide some of the features of ATM's VCs (virtual circuits) to an IP-centric world. Members of the IETF (Internet Engineering Task Force) who worked on MPLS recognized that ATM's advantages in service guarantees and VC visibility were too important to abandon in a TCP/IP-centric world and needed to be available in a connectionless domain.

The motivation for the MPLS proposal at that time was the latency in performing full IP longest-prefix matches at each hop in a routed network. By attaching an MPLS shim header to a packet at ingress, the network could identify packets and forward them in a switched network with speeds that could easily achieve SONET/SDH-like resiliency targets of less than 50 msec. Since MPLS first became a standard, silicon advances have allowed routed networks without MPLS to achieve high forwarding speeds. Nevertheless, both MPLS and PBB-TE have achieved justification outside performance parameters by offering a means of creating point-to-point "tunnels" that provide an easy means of offering services such as pseudowire and VPNs (virtual private networks).

The full G.8110 MPLS allows but does not mandate the use of routing and signaling protocols for topology and resource discovery. A centralized network-management system can fully provision MPLS LSPs (label-switched paths) without the use of any localized routing or signaling protocols. To further simplify the use of MPLS, the IETF identified the MPLS-TP (transport-profile) subset, which assumes the use of a connection-oriented packet transport employing a centralized static provisioning model, eliminating all local routing and signaling protocols. This step eliminates many protocols that developers built for MPLS, such as ECMP (equal-cost multiple path) and LSP merge. The MPLS-TP model allows the use of multiple client layers, including MPLS, native Ethernet, and TDM (time-division multiplexing).

PBB-TE and its predecessor, PBB, grew out of IEEE 802.1ad's QinQ effort, which allows service providers to create VLANs (virtual LANs) with a carrier-provided tag. Although a VPN employing MPLS hides a customer's MAC address, the original QinQ did not hide this address, requiring the entire network core to learn clients' MAC addresses. PBB in its original form, 802.1ah, dispensed with the spanning-tree protocols that 802.1ad used and eliminated the broadcasting of unknown MACs. Because of the way it defined MACs, its new header, MAC-in-MAC, referred to the complete re-encapsulation of a customer frame at the edge of the public network, hiding the customer address from the service-provider core.

Although both PBB-TE and MPLS can offer VPLS (virtual-private-LAN service), PBB offers these services as a bridged Ethernet tunnel. MPLS offers protected LSPs that users can centrally engineer and protect. The IEEE's incentive to move to 802.1Qay, or PBB-TE, was to overcome some of the perceived limitations of standard PBB. In the new PBB-TE standard, MAC learning and spanning-tree protocols are turned off for all backbone VLANs. Either a centralized network-management system or a GMPLS (generalized MPLS) control plane populates bridge tables with a static forwarding entry.

When the founders of the MEF (Metro Ethernet Forum) first met in 2001, their goal was to develop carrier service types, using strict Layer 2 systems when possible. Consequently, the MEF was closely associated with bridging and QinQ VLANs as it developed E-Line, E-LAN, and E-Tree (point-to-multipoint) service offerings (**Figure 1**). Although the MEF does not take an official position on pure Layer 2

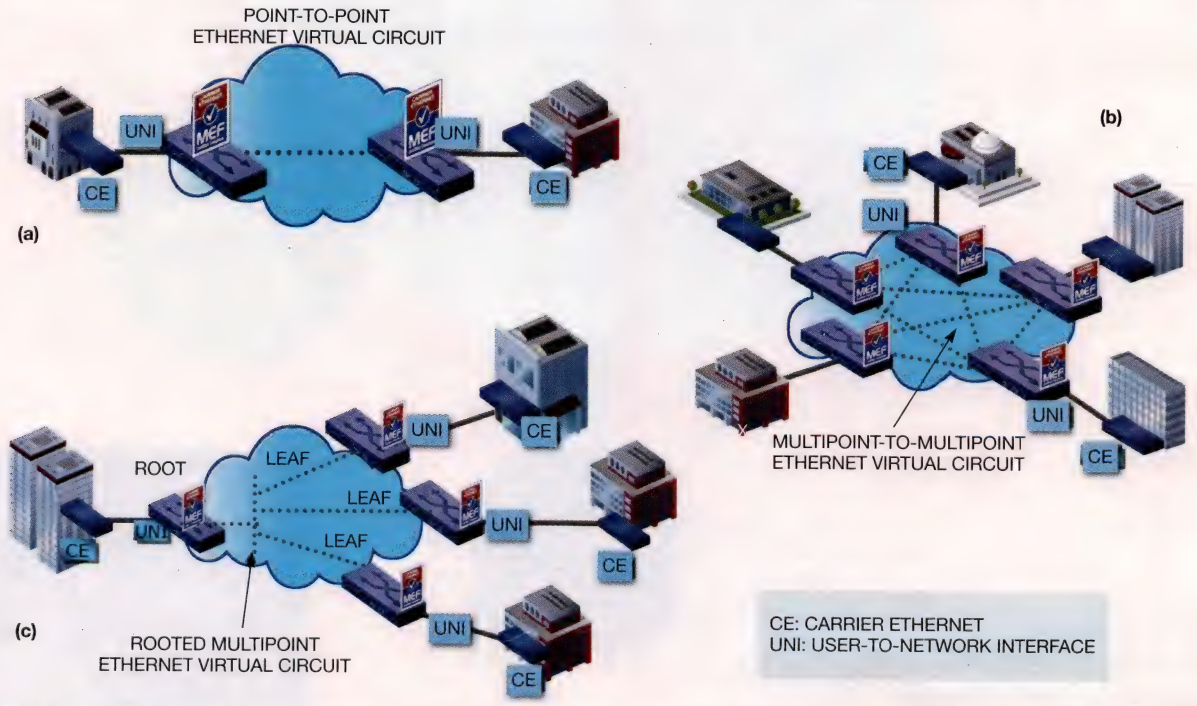

Figure 1 The Metro Ethernet Forum may have been thinking of bridge technology when it defined E-Line (a), E-LAN (b), and E-Tree (c) services, but MPLS can still be a better approach.

Thermal Challenge: VRM

High Ambient / High Watt Dissipation / Low Air Flow

Solution: New Bergquist S-Class Gap Pad® 5000S35

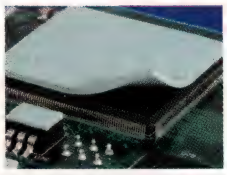

Gap Pad 5000S35's natural tack makes application clean and easy to handle.

Gap Pad S-Class is perfectly suited for high performance applications such as VRMs, BGAs and ASICs.

Bergquist's newest S-Class is the perfect combination of softness, low thermal resistance and high thermal conductivity. With a low bulk hardness (35 Shore 00) and high thermal conductivity (5.0 W/m-K) it conforms to demanding contours while maintaining its structural integrity. It is an ideal gap filling solution for applications with fragile components that can be damaged by harder materials that cause higher mounting pressure on components. Gap Pad 5000S35 is also an excellent solution for DVD drives, memory modules, and PC boards to chassis.

Excellent interfacing and wet-out makes GP 5000S35 ideal for fragile components with demanding contours and stack-up tolerances.

Ultra soft S-Class integrity offers easy application.

Gap Pad 5000S35 has a natural tack that eliminates the need for additional adhesive layers that can inhibit thermal performance. Its super soft, yet elastic nature provides excellent interfacing and wet-out, even to surfaces with high roughness or topography. Gap Pad 5000S35 features an embedded-fiberglass reinforcement that makes it puncture, shear and tear resistant. No tearing, flaking or crumbling – just clean and easy handling during the assembly process.

FREE S-Class swatch kit.

Visit our web site or call to qualify for your FREE S-Class swatch kit with product details.

FREE Gap Pad S-Class Swatch Kit

Call **1.800.347.4572** or visit **www.bergquistcompany.com/coolrun**

THE BERGQUIST COMPANY

18930 West 78th Street • Chanhassen, MN 55317 • A ISO9001:2000 registered facility
(800) 347-4572 • Phone (952) 835-2322 • Fax (952) 835-0430 • www.bergquistcompany.com

Thermal Materials • **Thermal Substrates** • **Fans and Blowers**

versus MPLS Layer 2.5 approaches and emphasizes the cost-effectiveness of pure bridging links for smaller networks, the forum optimized the service profiles for MPLS-based networks. Because the MEF advocates advanced OAM (operations/administration/maintenance) and performance-management tools, it supports the convergence of MPLS and emerging IEEE standards, such as 1588 (synchronous Ethernet) and 802.1ag (Ethernet fault management).

In theory, PBB-TE can offer most of the features of CE (carrier-Ethernet) service creation and tunneling, using a simpler hardware implementation employing Layer 2 switches and small per-port memories. PBB-TE is still at least one to two years from finalization, however. Even now, the cost delta is shrinking between PBB-TE and transport MPLS in multiport Ethernet chips. Does it continue to make sense to use PBB-TE in cheaper networks closer to the edge and preserve MPLS for metro core and aggregation duties, or does end-to-end MPLS make sense from an SOC (system-on-chip) implementation perspective in the near future?

MPLS has some unique functions that may be critical for some carriers. For example, it uses a label stack with a 3-bit traffic-class field that defines both the QOS (quality-of-service) level and explicit congestion notification. Careful implementation of 802.1Qay can offer some QOS features, but MPLS carries a natural advantage in defining classes of traffic for preferred delivery options in service-level agreements. Fast-reroute features in MPLS more easily provide the ability to hit SONET/SDH-like less-than-50-msec protection switching speeds than many Layer 2 alternatives can offer.

Market-analysis companies, including IDC (www.idc.com), have produced studies that indicate that upgrading Ethernet switches for PBB-TE support will add 33% to real estate and overall design costs. In contrast, the cost of adding Layer 3 functions for MPLS can vary tremendously based on factors such as numbers of ports. Therefore, no transport-tunneling method is cost-free. Because PBB-TE eventually will migrate to GMPLS as a control plane, the issue is where and why a low-cost bridging approach employing native Ethernet transport might make sense over comprehensive MPLS support.

Even if you assume that PBB-TE in its finalized form was technically antiquated upon completion of the standard, you can scarcely call the draft technology obsolete in the market of 2011. At the beginning of 2010, PacketExchange, an IP-specialized service provider using MPLS transport, acquired carrier Mzima for its PBB-TE expertise. Service providers standardizing on TCP/IP at layers 3 and 4 and Ethernet as a Layer 2 packet-framing method are coming to a common conclusion on the coexistence of MPLS and PBB-TE. As MPLS becomes more cost-effective to consider outside the network core, it may find more use than does PBB-TE. PBB-TE bridging technology, however, likely will have a place at the edge of the public network for years to come.

It makes sense for service providers to plan for MPLS-dominated networks. Thus, over time, Ethernet-chip developers must design for OEM products that emphasize greater MPLS features. The IEEE and MEF both see Layer 2 bridging as a constrained alternative to the multiple layers of transport guarantees that MPLS offers as a transport method. A mature Ethernet technology in metropolitan carrier networks should include such features as protection switching and distributed timing. In such networks, CE topologies can be natural adjuncts to the sort of IP/MPLS networking strategy that operates in both service-provider networks and large-scale enterprises. Silicon providers must accordingly offer comprehensive suites of switching, MAC, PHY, and coprocessor chips, with embedded support for PBB-TE and transport MPLS, allowing the OEM and, thus, carrier customers to carefully design the transport for the task at hand.

PBB-TE may be less strategic over time when planning for global backbone networks that large carriers manage. It can be ideal, however, for smaller regional carriers upgrading TDM-based metropolitan subnetworks or defining a new "greenfield"-network topology—that is, designing and implementing a network from the ground up. PBB-TE plays its most explicit role in metro networks largely employing full-chassis fault-resilient Ethernet switches. Today's Ethernet switches that seek to take on a carrier-transport role for a metropolitan-service-provider network must show fault tolerance, determinism, and sufficiently accurate network timing to allow transport of isochronous traffic.

Provided that a carrier network is well-characterized, PBB-TE can fulfill carrier-network QOS goals that have risen to a level that no one expected when developers recently designed some of the first VLAN-tagging methods. Networks can realize OAM functions, for example, through the use of the IEEE's 802.1ag standard for connectivity fault management. The network can establish a path-protection mechanism similar to that of SONET/SDH by creating separate backbone-VLAN identifiers for a work channel and a protection channel. With such extensions, the 10-Gbps CE switches now on the market look more like fault-resilient digital cross-connects than like enterprise Layer 2 switches.

Although PBB-TE can offer sophisticated OAM and performance-management functions, many emerging ITU (International Telecommunications Union) standards, such as Y.1731 and G.8031, work better in an environment extending MPLS to the network edge (**Figure 2**). A supporter of PBB-TE might insist that Y.1731's more expansive OAM goals depend more on new dedicated PHY

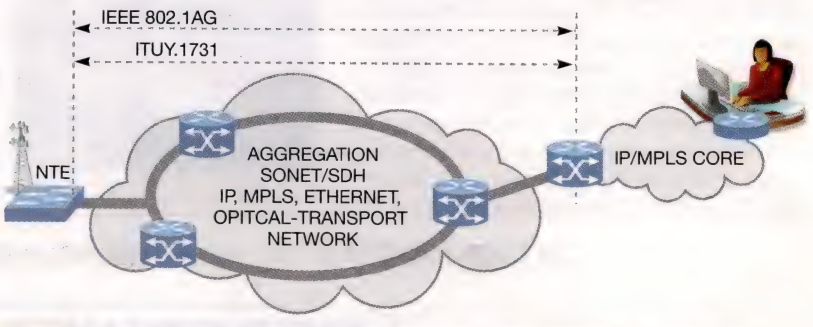

Figure 2 Network-management schemes can work best in an end-to-end MPLS environment.

EDN's Innovation Awards honor the most innovative technological advances in the electronics industry and the designers who created them. Join us for an evening of cocktails and hors d'oeuvres to introduce and celebrate the winners on May 2nd at the Corinthian Ballroom in San Jose, California.

Purchase tickets online now at http://innovation.edn.com/.

Award Categories Include:

- **Innovator of the Year:**
  Given to an electronics engineer or engineering team for outstanding innovation in product design or technical leadership
- **Innovations of the Year:**
  These awards honor unique, state-of-the-art electronics products in several categories

View the list of finalists and cast your votes for the winners at http://innovation.edn.com/ before March 31st.

Be there when the winners are announced!
Go to: http://innovation.edn.com/
Tickets are limited – Buy today!

EDN

SPONSORED BY AVNET Embedded

capabilities at the semiconductor or subsystem level than on a shift from bridging to Layer 2.5 label-switching. To a certain extent, that belief is true. Under MPLS, however, analysis of frame delay, locked conditions, packet loss, and the like can take place in an environment of broader traffic encapsulation and VC creation than is possible in a strict PBB-TE environment.

For example, G.8031 defines a signaling protocol that turns a VLAN link into an automatic-protection switching path, so designers could add the G.8031 approach to VLAN tags in duplicated backbone-VLAN identifiers that PBB-TE defines. The wealth of unidirectional and bidirectional protection-switching methods in G.8031 maps better into an MPLS topology, however, and both transport methods ultimately want to move to a GMPLS control plane, eliminating the differences between the two camps.

Recent advances in edge routing indicate that full Layer 3 control-plane approaches can move closer to the edge, although an access approach in which PBB-TE makes the most sense will also likely remain. The critical issue is whether MPLS serves strictly in a transport dimension or as a full control plane. The scalability of the control plane becomes the critical issue as MPLS moves to the access platform. The preprovisioning of physical or virtual circuits often makes more sense than extending MPLS control planes over large topologies.

THE CASE FOR END-TO-END MPLS TRANSPORT SEEMS TO HAVE A BROADER RATIONALE THAN THAT FOR PBB-TE.

Specialized metro service providers considering PBB-TE should expect equipment OEMs to upgrade simple bridged-Ethernet topologies in the public network with hardware and firmware supporting fault resiliency and protection-switching features, many of which 802.1Qay spells out. As a result, PBB-TE will continue for many years to come to play a role in implementing MEF services. It is too soon to bet against MPLS. As semiconductor features shrink and SOC designs become larger and more complex for a smaller unit price delta over simpler chips, the case for end-to-end MPLS transport seems to have a broader rationale than that for PBB-TE. EDN

AUTHOR'S BIOGRAPHY

Martin Nuss, PhD, is vice president, technology and strategy, for Vitesse and has more than 20 years of technical and management experience. Most recently, as CTO (chief technology officer) of Ciena's Optical Ethernet group, Nuss led the successful integration of the FlexSelect CN 4200 Metro WDM product line acquired from Internet Photonics, where he was founder and CTO. He also served 15 years at Bell Labs in various technical and management roles, including director of the optical-data-networks research department. Nuss is a fellow of the Optical Society of America and a member of the IEEE. He holds a doctorate in applied physics from the Technical University of Munich (Munich, Germany).

UBM Canon

Your Project Acceleration Resource for Electronics Manufacturing

Electronics® NEW ENGLAND

April 6–7, 2011

Boston Convention
& Exhibition Center
Boston, MA

Electronics-NE.com

Electronics® WEST

February 14–16, 2012

Anaheim Convention Center
Anaheim, CA

ElectronicsWestShow.com

NEW

Electronics® SOUTH

May 2–3, 2012

Charlotte Convention Center
Charlotte, NC

Find the suppliers, tools, and services you need to make your product, process, and business more efficient, cost-effective, and profitable.

For information on attending or exhibiting,
please call 310/445-4200 or visit CanonTradeShows.com

Produced and Managed by UBM Canon, 11444 W. Olympic Blvd., Los Angeles, CA, 90064-1549 Tel: 310/445-4200 UBMCanon.com

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Anticipator circuit speeds signal settling to a final value

Tim Regan, Linear Technology, San Jose, CA

The circuit in this Design Idea anticipates, or jumps instantly to, the final voltage of an input-signal change. It relies on the fact that the intended input signal changes exponentially with a known time constant. This circuit was adapted from a 1970s-era instant-reading electronic thermometer, which displayed a patient's body temperature within seconds after a thermometer probe was placed under the patient's tongue. It exploits the fact that the temperature probe's exponential thermal response-time characteristic is known.

The circuit uses a quad rail-to-rail amplifier to perform a mathematical operation (**Figure 1**). The input to the circuit is at Node X. At that node, a filter with a 500-msec RC time constant averages a 1-kHz PWM (pulse-width-modulated) signal. The desired output is a dc voltage proportional to the PWM duty cycle. A long time constant is required to reduce ripple. You obtain an instant output response by differentiating this input signal with the same time constant. The input signal is the voltage on capacitor C_1 as it moves from initial voltage V_I to final voltage V_F. R_1 and C_1 set the time constant, as the following **equation** shows:

$$V_{IN}=V_F-(V_F-V_I)\times e^{-\frac{t}{R_1 C_1}}=V_A,$$

where e is an irrational constant approximately equal to 2.718281828. You then buffer this signal with an inverting gain of one-half to prevent clipping. Ignoring dc biasing for clarity, the ac output at Node V_B is a function of the RC time constant, as the following **equation** shows:

$$V_B=-0.5V_F+0.5(V_F-V_I)\times e^{-\frac{t}{R_1 C_1}}.$$

You then differentiate the inverted signal with amplifier IC_1. You set the differentiator time constant with R_2 and C_2. The gain of a differentiator circuit increases with frequency, making these circuits prone to instability. You use R_5 and C_8 to keep the circuit stable. At the low frequencies of interest, R_2 and C_2 dictate the function of the circuit, as the following **equation** shows:

$$V_C=-\tau\times\frac{dV_B}{dt}=$$

$$-R_2C_2\times 0.5(V_F-V_I)\times e^{-\frac{t}{R_1 C_1}}\times\left(\frac{-1}{R_1C_1}\right).$$

R_1 and C_1 set the time constant of the input, so you can match it by making the differentiator time constant, $R_2\times C_2$, the same. This step cancels terms in the **equation** and simplifies the expression for output voltage, as the following **equation** shows:

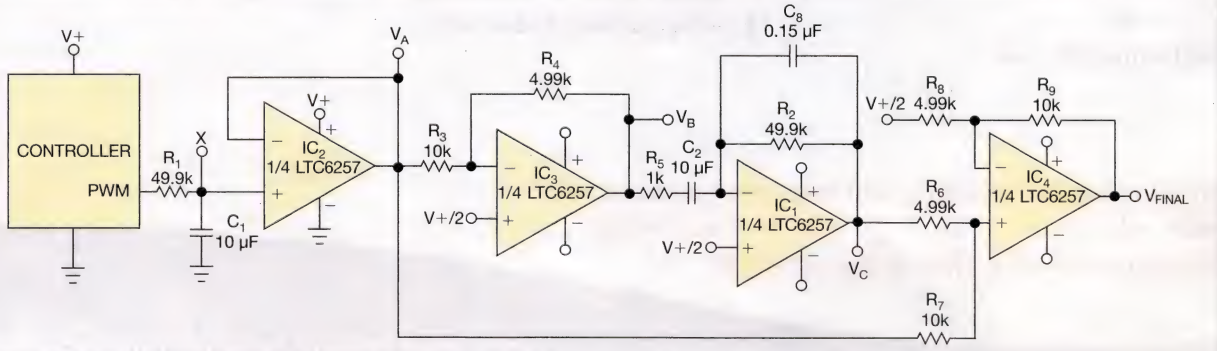

Figure 1 The output of this anticipator circuit instantly jumps to the final voltage of an exponential input signal.

DIs Inside

What are your design problems and solutions? Publish them here and receive $150! Send your Design Ideas to edndesignideas@ubm.com.

To see all of *EDN*'s Design Ideas, visit www.edn.com/designideas.

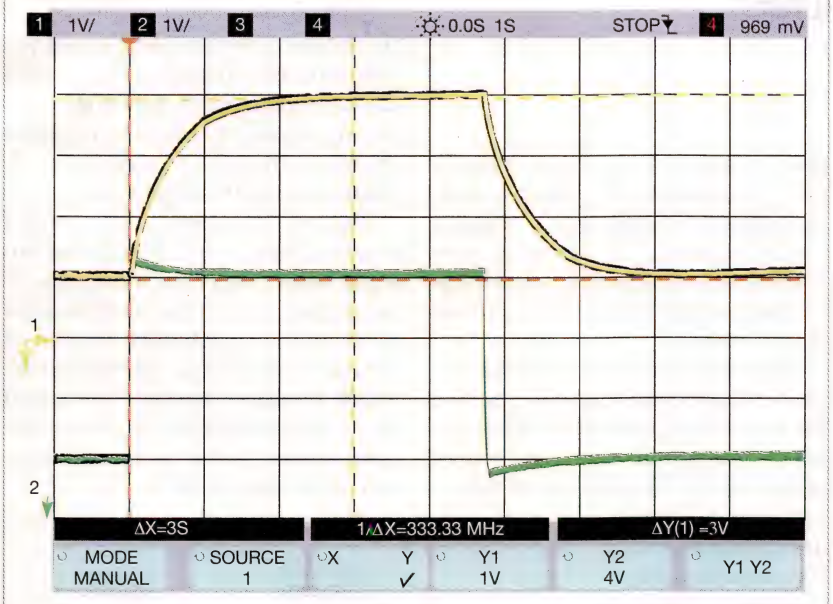

Figure 2 The anticipator circuit speeds the response of a slow exponential waveform (yellow) and results in nearly instant response to the final value, with only a small amount of overshoot (green).

$$V_C=0.5(V_F-V_I)\times e^{-\frac{t}{R_I C_I}}.$$

Due to scaling to prevent clipping, you sum this signal with the input signal in a weighted manner and present this voltage at the positive input of IC_4, as the following **equation** shows.

$$V_{IC4+IN}=\frac{2}{3}V_C+\frac{1}{3}V_A=\frac{(V_F-V_I)\times e^{-\frac{t}{R_I C_I}}}{3}+\frac{V_F}{3}-\frac{(V_F-V_I)\times e^{-\frac{t}{R_I C_I}}}{3}.$$

Note that the first and last terms of the preceding **equation** cancel out. You then set a gain of three for amplifier IC_4, as the following **equation** shows:

$$V_{OUT}=3\times\frac{V_F}{3}=V_{FINAL}.$$

When the input starts to move with a known exponential rate, the output anticipates the result and jumps instantly to what will be the final voltage (**Figure 2**). You can use this circuit in many applications that have a fixed input time constant. EDN

Schmitt trigger provides toggle function

Vladimir Oleynik, Moscow, Russia

D-type flip-flops, with the inverted output connected to the D input, can toggle a clock signal. If your circuit has an extra Schmitt-trigger inverter gate, you can use it to accomplish the same thing. Every time you press and release the momentary pushbutton switch in **Figure 1**, the circuit reverses to the opposite of what it was before you pressed it.

You can use a Schmitt trigger because it has hysteresis that positions roughly symmetrically around half of the power-supply voltage, $V_{CC}/2$. If a signal rises from 0V to $V_{CC}/2$, the output will be logic zero, so the inverter output is at logic one. If the signal drops from V_{CC} to $V_{CC}/2$, it will be at logic one, so the inverter output is at logic zero. At power-up, voltage at Point A rises from 0V to $V_{CC}/2$ because R_1=R_2, and the inverter's output will be high as V_{CC}. Capacitor C charges to $V_{CC}/2$. The momentary pushbutton switch in the circuit has NO (normally open) and NC (normally closed) states.

When you first press the button, capacitor C quickly charges to V_{CC} because the inverter output is high. When you release the button, high logic voltage appears across the capacitor on the inverter input. Thus, its output goes low. The capacitor discharges through resistor R_2 down to $V_{CC}/2$. When you press the button for the second time, capacitor C quickly discharges from $V_{CC}/2$ to 0V because the inverter's output is low. When you release the button again, the discharged capacitor shunts the inverter input, thus forcing its output high. The capacitor charges up through resistor R_1 to $V_{CC}/2$. **Figure 2** shows the waveforms.

The circuit is insensitive to contact bounce. Because the 40106 contains six Schmitt-trigger inverters, one IC can support as many as six momentary switches. You can substitute a two-input NAND Schmitt trigger CD4093 for the 40106. If you need to change the output when you press the button, reverse the connection of the NO and the NC contacts of the switch. EDN

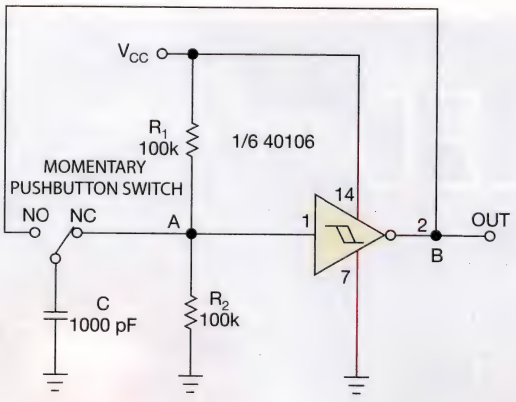

Figure 1 A single Schmitt trigger provides a toggle function.

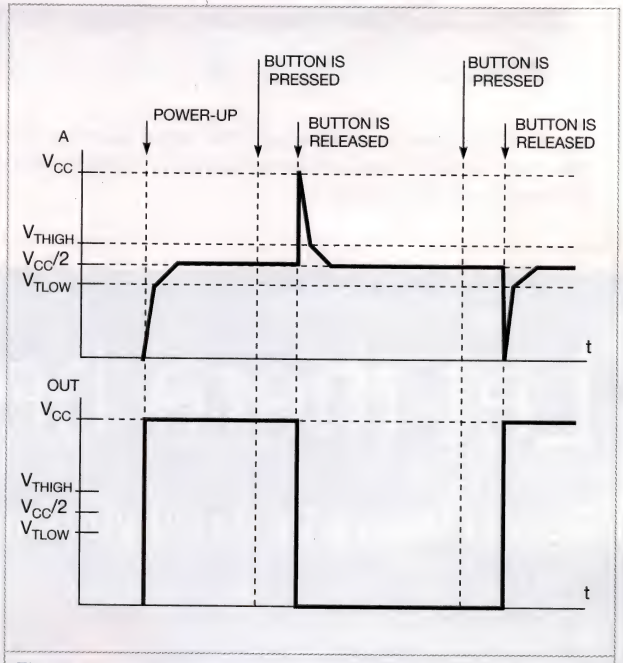

Figure 2 A waveform shows the signal plot in different points of the circuit.

Active multiplexing saves inputs

JB Guiot, Mulhouse, France

Microcontrollers must often read the status of switches in control applications. A typical switch configuration uses pullup resistors on both of the switches to pull the signals high or low for the microcontroller to read. The controlling output in the circuit in **Figure 1** uses two switches. When both switches are open, resistors R_1 and R_2 keep the input at an undetermined value between low and high, which is 1.5V if the supply voltage is 5V. Thus, an analog input would be preferable to a digital input. With the values in the figure, 0.5 mA of current flows through the resistors, even when both switches

AN INVERTER GATE HAS CLEARLY DEFINED LEVELS OF INPUTS. THUS, YOU CAN USE NORMAL LOGIC INPUTS RATHER THAN ANALOG INPUTS.

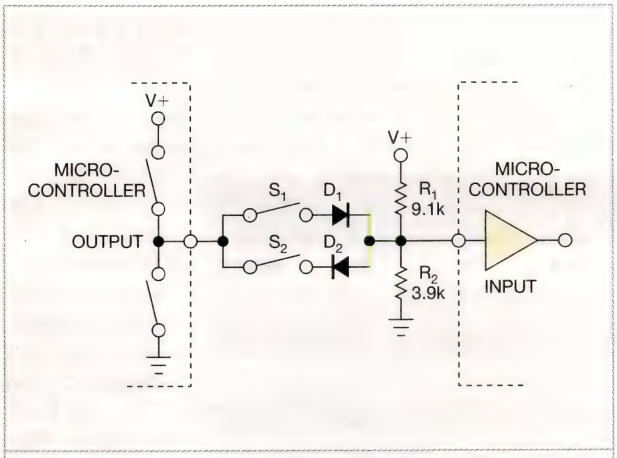

Figure 1 The controlling output in this circuit uses two switches.

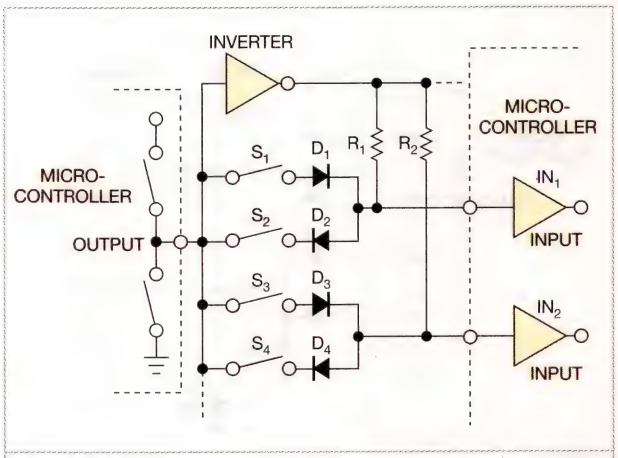

Figure 2 An inverter gate pulls the microcontroller's inputs to the opposite state of the driving output.

are open. You must multiply that value by the number of inputs used to get the total current.

Figure 2 shows an alternative circuit. It adds an inverter gate that pulls the microcontroller inputs to the opposite state of the driving output. An inverter gate has clearly defined levels of inputs. Thus, you can use normal logic inputs rather than analog inputs. You can choose a resistor value as high as the input characteristics allow yet low enough to minimize noise immunity. When both switches are open or the controlling output is in a high-impedance state, the current flowing through the resistor is only the current that flows into the microcontroller's input.

Switching a signal's state uses energy, so change states only when reading the switches. Otherwise, leave the input and output pins in a high-impedance state. If your design has position-cam switches that never close

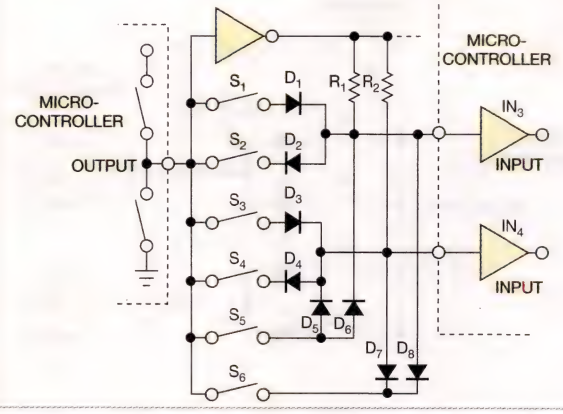

Figure 3 You can use this circuit to reduce the number of inputs coding the switches.

simultaneously, you can refer to the circuit in **Figure 3** to reduce the number of inputs coding the switches. **Table 1**, a truth table, provides the possible states.

You can expand this scheme to any number of switches and inputs. You can read 12 inputs with three switches. You can also mix the circuits in **figures 2** and **3** on the same microcontroller, separating independent switches (**Figure 2**) and "interlocked" switches (**Figure 3**) on different inputs. EDN

TABLE 1 TRUTH TABLE OF POSSIBLE STATES

Out	In_3	In_4	Switch
1	1	0	1
0	0	1	2
1	0	1	3
0	1	0	4
1	1	1	5
0	0	0	6

Transistor tester identifies terminals

Raju R Baddi, Tata Institute of Fundamental Research, Maharashtra, India

The simple transistor tester in **Figure 1** lets you identify the type of transistor, and it helps in detecting a transistor's emitter, collector, and base. It checks all possible combinations of directions of current flow between the test transistor's three terminals, T_1, T_2, and T_3.

The circuit uses two CD4022 or CD4017 counters, IC_1 and IC_2; a single-gate square-wave oscillator, G_4; and a CD4011 quad NAND gate, G_1 through G_3. A pair of LEDs connects in series to each test terminal to indicate the direction of current flow. The color of the LEDs directly reveals the junction side of the transistor.

Figure 2 provides an easy reference for understanding the test procedure. A pair of NPN transistors, Q_1 and Q_3, and PNP transistors, Q_4 and Q_6, for each terminal connects the terminals to either –V or +V, which sets up the required potential difference between the terminals. The circuit generates all of the possible or required combinations of +V and –V between the terminals to establish the junction relations. Q_7 and Q_8 act as voltage

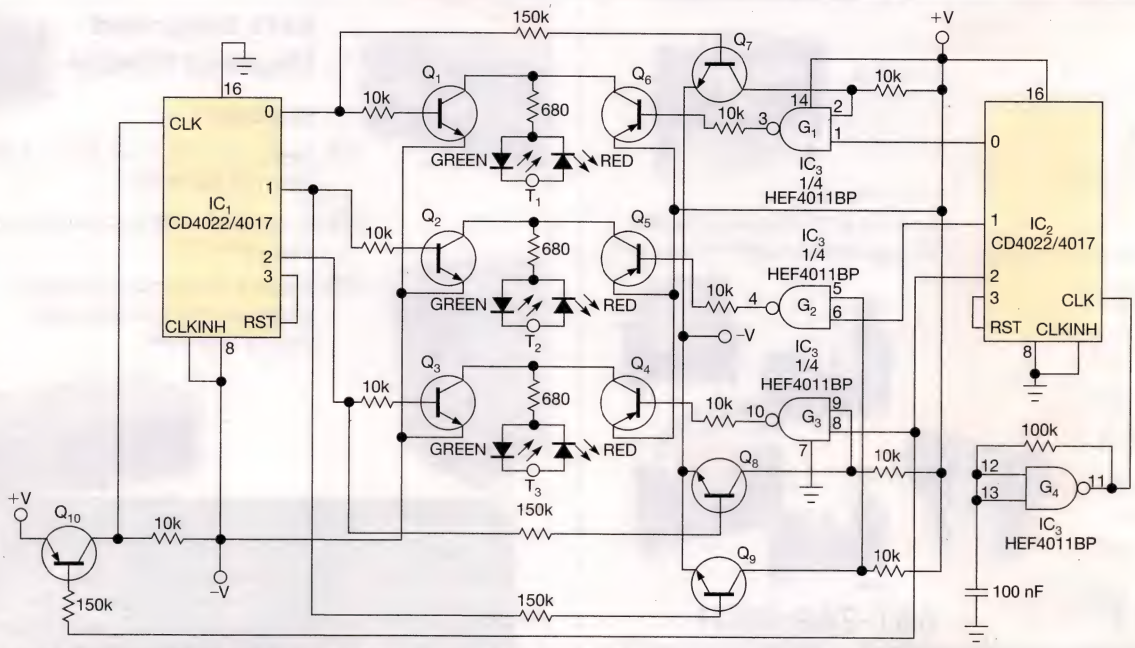

NOTE: NPN TRANSISTORS ARE BC547s, AND PNP TRANSISTORS ARE BC557s.

Figure 1 Counters and transistors let you detect a transistor's terminals and determine whether the transistor under test is working.

Introducing EDN's New Microsite on LEDs

Your complete source for up-to-the-minute information on LEDs and their applications.

- Latest News from the Editors at EDN, Design News and other Engineering Magazines
- Blogs, including Margery Conner's Popular "Power-Source" Blog
- Resource Center, Including White Papers, eBooks, Tutorials, Webcasts and More
- Videos, Including Design Tips, Application Demos, and More
- Updates on Future In-Person Workshops and Seminars, like this one!
- Fun Challenges to Test Your Knowledge of LEDs

Check It Out Today at www.edn.com/leds

translators, whereas G_1 to G_3 are inhibitors, which prevent T_1 to T_3 from clashing by being at +V and –V at the same time.

When you plug a functioning transistor into the test terminals, it restricts current flow in certain directions only. The series LEDs reveal these directions and, hence, indicate the type of transistor—for example, the LEDs glow red-green-red for an NPN transistor and glow green-red-green for a PNP transistor.

With this knowledge, you can easily choose the base of the transistor. To differentiate between emitter and collector you must understand the property that, under reverse bias, base-emitter junction breaks down more easily than does the base-collector junction, which is reverse-biased for normal operation.

Because transistors have different base-emitter reverse-breakdown voltages, the circuit provides a way to easily change the supply voltage (**Figure 3**). Under increased voltage, both LEDs connected with the emitter glow brightly, whereas only one LED glows for the collector (**Figure 2b** and **d**). A basic voltage of ±4V seems sufficient for detecting the base or type of transistor. By gradually increasing the supply voltage from ±4 to ±15V, you can test a variety of transistors for the emitter. This range provides a maximum reverse-breakdown voltage of greater than 26V for the base-emitter junction, taking into account the voltage drop of the series LEDs.

This circuit underwent testing and works. However, the testing employed CD4520 counters and CD4028 decoders because the CD4022/CD4017 ICs were unavailable. This replacement shouldn't cause problems. Only the voltage levels matter, which for CMOS devices is more or less the same for logic one or logic zero. You can also use only two supply voltages: ±5V for detecting the base and ±15V for detecting the emitter. EDN

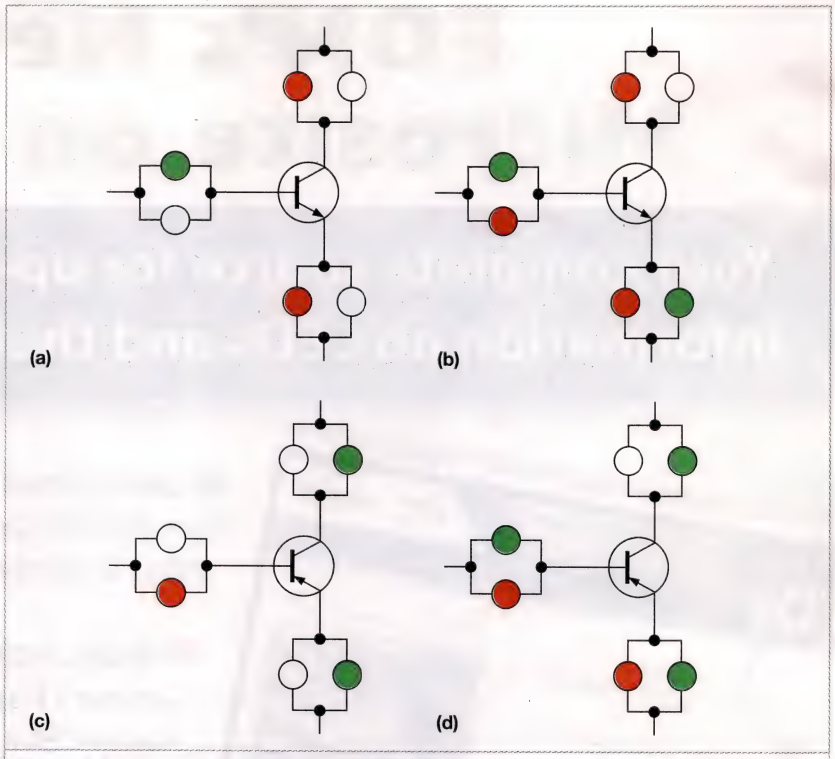

Figure 2 A reference indicates the health of a transistor under test for an NPN with normal supply voltage (a), an NPN with increased supply voltage (b), a PNP with normal supply voltage (c), and a PNP with increased supply voltage (d).

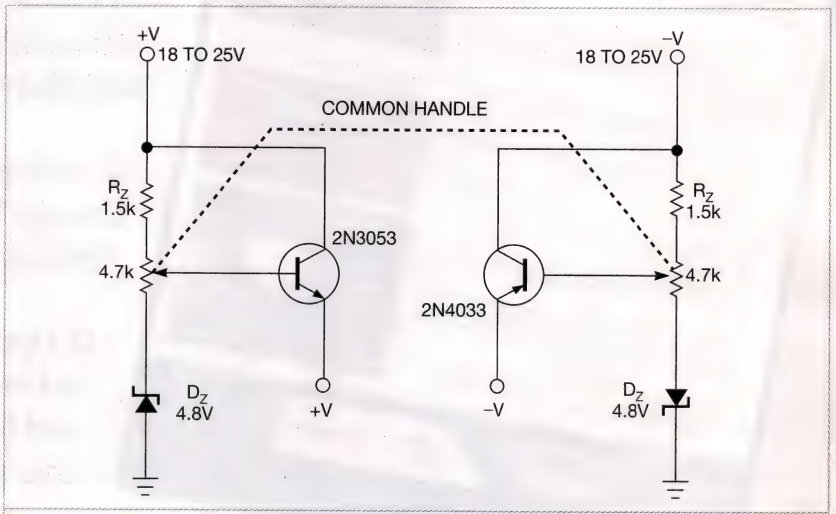

Figure 3 A common handle ensures symmetric variation of supply voltage about the ground.

Finely tune the hue of blue-light sources

Marián Štofka, Slovak University of Technology, Bratislava, Slovakia

When coach cars of long-range trains comprised compartments for six to eight passengers, the passengers could choose either "white" or deep-blue light. The blue light helped passengers sleep, even when they were not in full darkness. The circuit in **Figure 1** lets you set a hue of blue light that can match your favorite blue color. Although you can accomplish this task

with another circuit (**Reference 1**), this circuit provides a finer adjustment of color and uses a narrower range from cyan to royal blue. The light is a mixture from two power LEDs, both from Avago Technologies (www.avagotech.com). LED_1, the ASMT-JC11, is a recently introduced, high-performance, thermally ruggedized, miniature, 1W cyan unit, and LED_2 is the 3W ASMT-JL31.

Two coupled current sources drive the LEDs. IC_{3A}, with a cascade of two NPN transistors, Q_1 and Q_2, forms a sink-current source that drives LED_2. IC_{3B}, a current source with a cascade of two PNP transistors, Q_3 and Q_4, drives LED_1. This PNP current source is feed-forward-operated and is controlled by the output current of the sink-current NPN source. The circuit achieves this task by routing the output sink current through feedback resistor R_{E2} of the PNP current source. If the output sink current is at full-scale, then the feedback signal for the PNP current source is also at maximum. Thus, the actual sourced current has a theoretical value of 0 mA. In contrast, if the sink current is 0 mA, then the source current reaches the full-scale value. Therefore, the sink and the source currents are complementary; their sum is a constant. The sum of the output currents is $I_O=(R_B \times I_T)/R_E$, where I_T is the value of the reference current that flows through IC_1, an Analog Devices (www.analog.com) AD590. IC_1 is a two-pole source of proportional-to-absolute-temperature current, whose value is typically 298.2 µA at room temperature. IC_1 creates the high- and low-side reference voltages, V_{REF}, which are both 400 mV and which serve as references for the two power-current sources. I_O has a value of approximately 80 mA.

With values for R_{E2} and R_{E1} tightly matched and equal to R_E, the sums of the currents flowing through the LEDs are independent upon setting IC_2, an Analog Devices AD5228, which acts as a digitally controlled potentiometer. The sum of emitted light from the two LEDs remains roughly constant as they change the final hue from cyan to royal blue.

You can interchange the position of LEDs in the circuit, but using them as the **figure** shows offers optimum voltage headroom for both power-current sources, even though the forward-voltage drop of the cyan LED is higher. The wiper positions of the DAC have margins that are typically 0.9% for zero and −2.4% for full-scale. With the Preset input high, you set the midscale setting of the wiper, W, at power-up; thus, the output light is a 50/50 mixture of both colors. For a low Preset, you reach a zero setting, resulting in full-cyan light at power-on. IC_2's internal 100-kΩ resistors force the $\overline{PU}$ and $\overline{PD}$ control pins to inactive high. As I_T's value rises linearly with absolute temperature, the circuit roughly compensates the decreasing of radiance of the LEDs.

In advertising or toys, this circuit can also provide a periodic change of the hue. If you set the Preset high and hold the $\overline{PU}$ pin low while feeding a 50%-duty-cycle, 0.05-Hz-frequency logic waveform to the $\overline{PD}$ pin, you get a slow, periodic, quasicontinuous "waving" of the color from cyan to royal blue and back. EDN

REFERENCE

1 Štofka, Marián, "Electronically tinge white-light source," *EDN*, Nov 4, 2010, pg 46, http://bit.ly/eQhzR5.

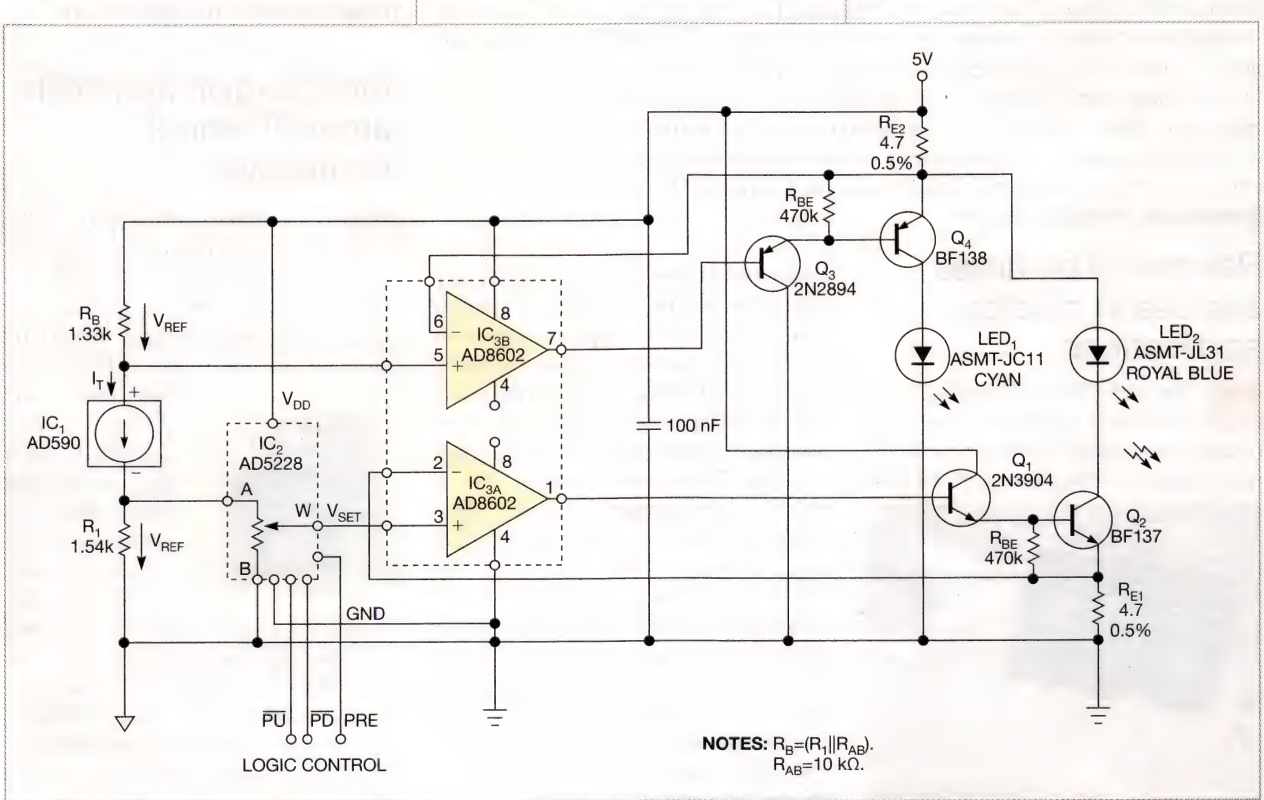

Figure 1 A current source and a current sink drive two LEDs. The balance of current sets the hue.

productroundup

AMPLIFIERS, OSCILLATORS, AND MIXERS

Op amps target audio applications

The dual-channel APX4558 op amp features input noise of 8 nV/√Hz at a 1-kHz unity-gain bandwidth of 3 MHz. Targeting use in cost-sensitive audio products, including set-top boxes and preamplifiers, the op amp comes in a commercial grade with an operating-temperature range of 0 to 70°C; an industrial-grade version, the APX4558I, operates from –40 to +105°C. Both parts come in SO-8 packages and have short-circuit protection. Power-supply voltage is ±18V, maximum offset voltage is 6 mV, and maximum offset voltage over temperature is 7.5 mV. Maximum input bias current is 500 nA, and maximum input bias current over temperature is 800 nA. The devices cost 8 cents (1000) each.

Diodes Inc, www.diodes.com

Rail-to-rail op amps suit use in medical applications

The MAX9636, MAX9637, and MAX9638 rail-to-rail low-noise operational amplifiers have 38-nV/√Hz input-voltage noise density, 50-fA √Hz input-current noise density, and 0.1-pA input bias current. The devices operate from a 2.1 to 5.5V supply and suit use in portable-medical-system applications. They feature a typical quiescent current of 36 µA and a low-power shutdown mode that reduces supply current to 1 µA. They operate over a –40 to +125°C temperature range. The MAX9636 comes in a lead-free, 2×2.2-mm, six-pin SC70 package; the MAX9637 comes in a lead-free, 2×2.2-mm, eight-pin SC70 package; and the MAX9638 comes in a lead-free, 1.4×1.8-mm, ten-pin UTQFN package. Prices start at 33, 50, and 60 cents (1000), respectively, for the MAX9636, MAX9637, and MAX9638.

Maxim Integrated Products, www.maxim-ic.com

Low-voltage op amps provide sensor measurement

The 1.8V OPA2320 op amp provides 20 MHz of gain bandwidth and 0.2-pA input bias current for high-impedance optical and sensor front ends, programmable-logic controllers, motor-control loops, and test-and-measurement equipment. Maximum offset voltage is 150 µV, and CMRR is 114 dB. The op amp has rail-to-rail I/O and noise of 7 nV/√Hz at 10 kHz and a quiescent current of 1.45 mA per channel. The device comes in 3×4.9-mm MSOP-8, 4.9×6-mm SOIC-8, and 3×3-mm SON-8 packages and sells for $1.25 (1000).

Texas Instruments, www.ti.com

Variable-gain amplifiers target RF and IF frequencies

The digitally controlled ADL5243 variable-gain amplifier operates in the 100- to 4000-MHz frequency range and employs GaAs technology. It combines a digital-step attenuator, a gain block, and a broadband ¼W driver amplifier. Gain control is 31.5 dB in 0.5-dB steps, and the device has parallel- and serial-interface modes. Linearity is greater than 40 dBm, and the noise figure is 2.9 at 900 MHz. The ADL5243's driver offers performance of 42 dBm with a 25.7-dBm 1-dB gain-compression measurement of 2.14 GHz. The device comes in a 5×5-mm, 32-lead LFCSP and sells for $13.98 (1000).

Analog Devices, www.analog.com

16V op amps feature 135-µA quiescent current

The MCP6H01 single and MCP6H02 dual operational amplifiers have a gain-bandwidth product of 1.2 MHz and a supply voltage of 3.5 to 16V. The devices feature a typical quiescent current of 135 µA, a maximum offset voltage of 3.5 mV, a typical common-mode rejection ratio of 100 dB, and a typical power-supply rejection ratio of 102 dB. The devices target use in medical applications, such as portable instrumentation and heart-and blood-pressure

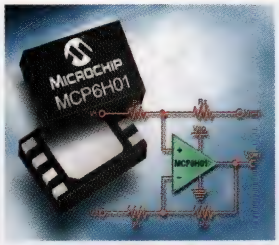

EDN ADVERTISER INDEX

Company	Page	Company	Page
Agilent Technologies	31, 41	MathWorks Inc	27
Allied Electronics Inc	25	Memory Protection Devices	61
Analog Devices Inc	23	Micrel Semiconductor	43
austriamicrosystems AG	59	Micro Crystal	46
Avnet Electronics Marketing	10	Mornsun America LLC	33
Centellax	29	Mouser Electronics	6
Cirrus Logic Inc	35	National Instruments	13, 15, 17
Coilcraft	4	PalPilot	61
Dell	21	Panasonic Industrial	67
Digi-Key Corp	C-1, C-2	Pico Electronics Inc	37, 48, 49
Everlight Electronics Co Ltd	11	Renesas Technology Corp	51
Express PCB	46	RFM	47
Fox Electronics	67	Silicon Labs	C-3
Front Panel Express	67	Tektronix	3
Integrated Device Technology	8	The Bergquist Company	53
International Rectifier Corp	7	Trilogy Design	67
Lambda Americas	44	Ultraview Corp	67
Lattice Semiconductor	19	Xilinx	45
Linear Technology	C-4		

EDN provides this index as an additional service. The publisher assumes no liability for errors or omissions.

UBM
Canon

Accelerate Your Projects into the Future

AEROCON

AEROSPACE • DEFENSE • INSTRUMENTATION

MARCH 16-17, 2011
Orange County Convention Center
Orlando, Florida

SEPTEMBER 20-22, 2011
McCormick Place North
Chicago, IL

FEBRUARY 14-16, 2012
Anaheim Convention Center
Anaheim, California

MARCH 14-15, 2012
Fort Worth Convention Center
Fort Worth, Texas

Introducing the conference and exhibition for the aerospace and defense industries

Featuring suppliers for the design and production of aerospace and defense products, including:

- Software
- Machinery
- Materials
- Components
- Sub-Assemblies
- Contract Services
- Electronics
- Testing

...and much more!

Your Project Acceleration RESOURCE

For more information on attending or exhibiting, visit:

AeroConShows.com

Promo Code: AA

Produced and managed by UBM Canon | canontradeshows.com

monitors; automotive applications, such as proximity, temperature, or flow sensors; and industrial applications, such as high-side current sensing in power supplies. PCB footprints and schematic symbols are available at the vendor's Web site at www.microchip.com/get/2KPK. The downloads are available in a neutral format that you can export to EDA CAD/CAE design tools using the Ultra Librarian Reader from Accelerated Designs Inc (www.accelerated-designs.com). The MCP6H01-E/SN and MCP6H01T-E/SN come in eight-pin SOICs and sell for 47 cents (10,000) each, and the MCP6H01T-E/MNY comes in an eight-pin, 2×3-mm TDFN package and sells for 47 cents. The MCP6H02-E/SN and MCP6H02T-E/SN sell for 66 cents each and come in eight-pin SOICs, and the MCP-6H02T-E/MNY comes in a 2×3-mm TDFN package and sells for 66 cents.

Microchip Technology,
www.microchip.com

Audio power amplifier targets vehicle-entertainment equipment

The TDA7850LV audio power amplifier allows in-car entertainment equipment to operate without interruptions and distortions as the engine automatically turns off and restarts to conserve energy. The Class AB MOSFET amplifier has four channels providing 30W each into 4Ω with 14.4V power and 10% distortion at 1 kHz. It also has four channels providing 53W each into 2Ω with 14.4V power and 10% distortion at 1 kHz. The amplifier features automatic mute at minimum supply-voltage detection; 26-dB internally fixed gain; protection against very inductive loads, output short circuits to ground and to the supply voltage across the load, load-dump voltage, and reversed battery; and protection for chip temperature with a soft thermal limiter. It sells for $8.047 (1000).

STMicroelectronics,
www.st.com

Acid test

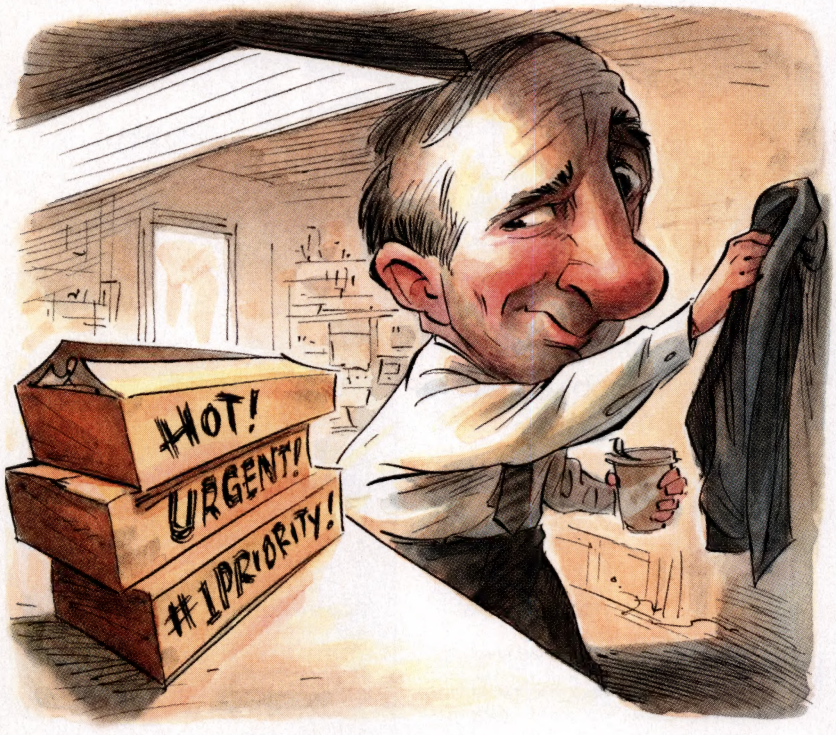

You can expect a few in-warranty failures in any complex product. However, when there's a sudden increase of such claims after the release of a new product, alarm bells should ring, and the design-engineering department should be involved. Such was the case in my company with a UV (ultraviolet)-light source and sensor assembly we were using to detect counterfeit currency. Counterfeits sometimes fluoresce when you expose them to UV light, whereas genuine US currency usually does not.

We had used this technique for many years in my company's products, but some months after a redesign to make the subassembly smaller and cheaper, an alert customer-service manager noticed that there were many more warranty claims than before. Within a day, three defective units found their way to my desk. With an excess of playfulness, my boss had marked them "Hot!" "Urgent!" and "#1 Priority!"

The redesigned units were black-plastic housings containing two CCFLs (cold-cathode fluorescent lamps), a photodiode, and a glass window to let the UV light out and any light resulting from fluorescence in. I worked with a mechanical engineer on the job, selecting UV-stable plastics, poring over the published transmission curves of filter glasses, and finding the best photodiode our budget could afford. Two halves of the units were screwed together, and, because the working environment was a currency-counting machine, they were sealed against paper-dust intrusion with an RTV (room-temperature-vulcanizing) adhesive.

As a first step in isolating the cause of these field failures, I plugged each unit into a test bed and saw that, in each case, the fluorescent lamps failed to light. I disassembled one of the plastic housings—the one marked "Hot!"—removing the screws and cutting through the adhesive with a sharp knife. A little inspection revealed that one of the bare copper wires coming from the UV lamps had broken. Copper isn't a brittle metal, and it seemed unusual for a wire to break unless someone had subjected it to repeated movement—an unlikely scenario inside a sealed unit. I thought that perhaps the wire had gotten nicked during assembly and was on the verge of breaking.

I continued on to the second unit—"Urgent!"—and quickly found that this one had not just one but two broken wires. With a magnifying glass, I could see that the broken ends were not clean cuts but were corroded. Corrosion is a chemical process, not electronic, so I was more than a bit out of my field. I bravely forged ahead, turning my thoughts toward the materials used in the unit: some sort of injection-molded plastic, fluorescent lamps, a blue-glass window, copper wires, RTV adhesive, plated-steel screws, and more.

A close examination of the "#1 Priority!" unit showed tiny droplets of moisture, just a thin film of fog, on the inside of the glass window, in addition to the now-expected broken wire. Could water vapor have corroded the copper wires? How could water vapor have entered a sealed unit? Could it be condensation of air humidity? Was this film really plain water? A little research revealed that, as some RTV adhesives cure, they release acetic acid, the same stuff that's in vinegar. And acetic acid corrodes copper. In this case, the adhesive and wires were trapped together in a sealed environment, making the corrosion even worse. After some consultation, the adhesive vendor recommended an electronics-grade RTV, which would cure without releasing acetic acid.

After a free retrofit of all units in the field, the in-warranty failure rate returned to zero, customers were happy, and I went on with a new appreciation for chemistry and a new dedication to multidisciplinary design reviews. EDN

David R Bryce currently works for SEPTA (Southeastern Pennsylvania Transportation Authority, Philadelphia). He lives with his beloved wife and useless cat in Morrisville, PA.

DANIEL VASCONCELLOS

SIMPLIFY THE ADDITION OF USB CONNECTIVITY IN EMBEDDED DESIGNS

Silicon Labs' new CP21xx USB bridge ICs eliminate software complexity and driver compatibility issues, providing a cost-effective, small-footprint turnkey solution for adding USB and human interface device (HID)-USB connectivity to applications such as personal medical devices, cellular and cordless phones, smart card and flash card readers, personal digital assistants, MP3 players, bar code readers, wireless modems and industrial control systems.

With the rapid proliferation of USB and HID-USB in the embedded world, developers are looking for painless, economical ways to embed USB connectivity in their designs. A leading supplier of USB connectivity solutions, Silicon Labs developed the CP21xx bridge family so that designers can add USB to microcontroller-based applications without the cost and complexity of developing sophisticated USB software. The CP21xx ICs require no USB expertise to deploy, enabling developers to focus time and resources on end applications.

USB-TO-UART BRIDGE BLOCK DIAGRAM

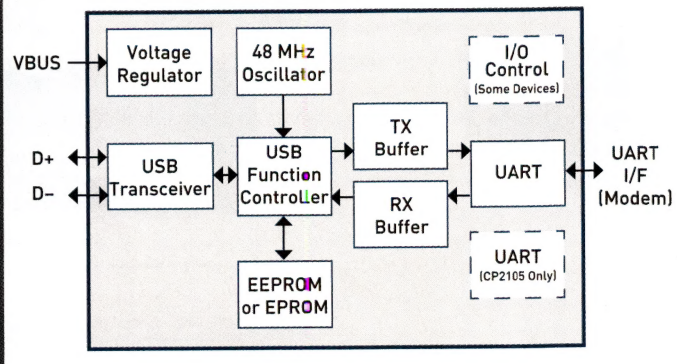

CP21xx FEATURES

- USB 2.0 compliant, full-speed (12 Mbps)
- No external crystal required
- Up to 1024 Bytes of EEPROM or EPROM
- User programmable custom Baud rates
- Supports all modem interface signals
- Baud rates up to 2 Mbps
- Industrial temperature range -40 to 85 °C
- SMBus master device (CP2112)
- HID class support, no drivers needed (CP2112)
- 512 B SMBus data buffer (CP2112)

APPLICATIONS

- USB to RS-232 converters
- USB to Dual RS-232 converters
- USB to RS-422/RS-285 converters
- Upgrade of legacy RS-232 devices
- PDA USB interface cable
- Cellular phone USB interface cable
- Barcode readers in Point-of-Sale terminals
- PC peripherals
- Industrial connectivity (CP2112)

EASY USB-TO-UART/I²C BRIDGE COMMUNICATION

See how easy it is to add USB to your embedded design:
Download the white paper "USB Simplified."
www.silabs.com/usb

SILICON LABS

36V, 4A Sync Buck

Actual Schematic

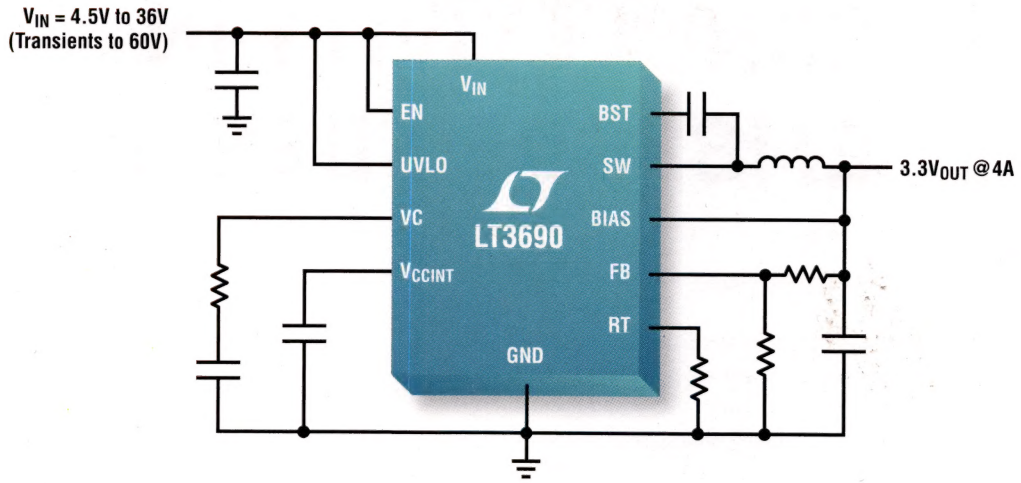

70µA I_Q, <15mV$_{P-P}$ Output Ripple and 92% Efficient

The LT®3690 brings a new level of performance and features for high voltage point-of-load step-down conversion. Its wide input voltage range of 3.9V to 36V, with 60V transients, meets the needs of many of the rails common to industrial, medical and automotive applications. Up to 4A of continuous output current can be supplied to the load with minimal thermal design. Its 70µA of quiescent current in standby maximizes battery life in "always-on" applications. With up to 1.5MHz switching frequency and a high level of integration, the external components are few and small, enabling a compact solution footprint.

Features

- Input Voltage Range: 3.9V to 36V (60V Transients)
- 4A of Continuous Output Current
- Quiescent Current of 70µA (12V$_{IN}$ to 3.3V$_{OUT}$)
- Low Output Ripple of <15mV$_{P-P}$ in Burst Mode® Operation
- Programmable Switching Frequency: 170kHz to 1.5MHz
- Output Voltage Range: 0.8V to 20V
- Programmable Input Undervoltage Lockout
- Compact 4mm x 6mm QFN Package

LT3690 Demo Circuit
(Actual Size)

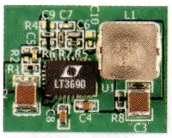

(20mm x 15mm)

Info & Free Samples

www.linear.com/3690

1-800-4-LINEAR

LINEAR TECHNOLOGY